

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

Lars E.O. Svensson
Sveriges Riksbank and Princeton University

Noah Williams
University of Wisconsin

Our previous work develops 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 (Svensson and Williams, 2007a, 2007b). 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 discuss how these modes could be structured to capture many different types of uncertainty relevant for policymakers. Here we put those suggestions into practice. We start by briefly discussing how an MJLQ model can be derived as a mode-dependent linear-quadratic approximation of an underlying nonlinear model, and we then apply our methods to a simple empirical mode-dependent New-Keynesian model of the U.S. economy, using a variant of a model by Lindé (2005).

In Svensson and Williams (2007b), we study optimal policy design in MJLQ models when policymakers can or cannot observe the current mode, but we abstract from any learning and inference about the

We thank James Bullard, Timothy Cogley, and Andrew Levin for comments on an earlier paper of ours, which helped inspire this paper, and Carl E. Walsh for comments on this paper. The views expressed in this paper are solely our responsibility 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.

Monetary Policy under Uncertainty and Learning, edited by Klaus Schmidt-Hebbel and Carl E. Walsh, Santiago, Chile. © 2009 Central Bank of Chile.

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 Svensson and Williams (2007a), we focus on learning and inference in the more relevant situation, particularly for model-uncertainty applications in which the modes are not directly observable. Thus, decisionmakers must filter their observations to make inferences about the current mode. As in most Bayesian learning problems, the optimal policy typically includes an experimentation component reflecting the endogeneity of information. This class of problems has a long history in economics, and solutions are notoriously difficult to obtain. We developed algorithms to solve numerically for the optimal policy.¹ Given 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 updates the probability distribution of the current mode in a Bayesian way, but the optimal policy is computed each period under the assumption that the policymaker will not learn from observations in the future. In our setting, the AOP is significantly easier to compute, and in many cases it 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 that have been widely studied in macroeconomics.³ The AOP specification also rules out the experimentation that some may view as objectionable in a policy context.⁴

In this paper, we apply our methodology to study optimal monetary policy design under uncertainty in dynamic stochastic

1. In addition to the classic literature (on such problems as a monopolist learning its demand curve), Wieland (2000, 2006) and Beck and Wieland (2002) examine Bayesian optimal policy and optimal experimentation in a context similar to ours but without forward-looking variables. Tesfaselassie, Schaling, and Eijffinger (2006) 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 (2001) and Cogley, Colacito, and Sargent (2007) study situations like ours, but their 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 (2006) analyzes active and passive learning in a New-Keynesian model with uncertainty about the slope of the Phillips curve.

2. The literature also refers to optimal policy under no learning, adaptive optimal policy, and Bayesian optimal policy as myopia, passive learning, and active learning, respectively.

3. See Evans and Honkapohja (2001) for an overview.

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

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 that at least partially reflects our assumption of symmetric information between the policymakers and the public—and then quantify the additional gains from experimentation.⁵

Since we typically find that the gains from experimentation are small, the rest of the paper focuses 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 that is a mode-dependent variant of Lindé (2005). 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.

The paper is organized as follows. Section 1 presents the MJLQ framework and summarizes our earlier work. Section 2 presents our analysis of learning and experimentation in a simple benchmark New-

5. In addition to our own previous work, MJLQ models have been widely studied in the control-theory literature for the special case in which the model modes are observable and there are no forward-looking variables (see Costa, Fragoso, and Marques, 2005, and the references therein). Do Val and Basar (1999) provide an application of an adaptive-control MJLQ problem in economics. Zampolli (2006) uses such an MJLQ model to examine monetary policy under shifts between regimes with and without an asset-market bubble. Blake and Zampolli (2006) extend 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 (2007b) 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 Marcat and Marimon (1998). That paper also provides two concrete examples: an estimated backward-looking model (a three-mode variant of Rudebusch and Svensson, 1999) and an estimated forward-looking model (a three-mode variant of Lindé, 2005). Svensson and Williams (2007b) 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 (2007a) focus on learning and experimentation in the MJLQ framework.

Keynesian model, and section 3 presents our analysis in an estimated empirical New-Keynesian model. Section 4 presents some conclusions and suggestions for further work.

1. MJLQ ANALYSIS OF OPTIMAL POLICY

This section summarizes our earlier work (Svensson and Williams, 2007a, 2007b). We start by describing our MJLQ model and then briefly discuss approximate MJLQ models. Finally, we explore the three types of optimal policies considered: optimal policy with no learning, adaptive optimal policy, and Bayesian optimal policy.

1.1 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 \mathbf{X}_t denote an n_X vector of predetermined variables in period t , \mathbf{x}_t an n_x vector of forward-looking variables, and \mathbf{i}_t an n_i vector of policymaker instruments (control variables).⁶ 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

$$\mathbf{X}_{t+1} = \mathbf{A}_{11j_{t+1}} \mathbf{X}_t + \mathbf{A}_{12j_{t+1}} \mathbf{x}_t + \mathbf{B}_{1j_{t+1}} \mathbf{i}_t + \mathbf{C}_{1j_{t+1}} \boldsymbol{\varepsilon}_{t+1}, \quad (1)$$

$$E_t \mathbf{H}_{j_{t+1}} \mathbf{x}_{t+1} = \mathbf{A}_{21j_t} \mathbf{X}_t + \mathbf{A}_{22j_t} \mathbf{x}_t + \mathbf{B}_{2j_t} \mathbf{i}_t + \mathbf{C}_{2j_t} \boldsymbol{\varepsilon}_t, \quad (2)$$

where $\boldsymbol{\varepsilon}_t$ is a multivariate normally distributed random i.i.d. n_ε vector of shocks with mean zero and contemporaneous covariance matrix $\mathbf{I}_{n_\varepsilon}$. The matrices \mathbf{A}_{11j} , \mathbf{A}_{12j} , ..., \mathbf{C}_{2j} have the appropriate dimensions and depend on the mode j . Given that 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.⁷

The matrices on the right-hand side of equation (1) depend on the mode j_{t+1} in period $t + 1$, whereas the matrices on the right-hand side

6. The first component of \mathbf{X}_t may be unity, to allow for mode-dependent intercepts in the model equations.

7. See also Svensson and Williams (2007b), where we show how many different types of uncertainty can be mapped into our MJLQ framework.

of equation (2) depend on the mode j_t in period t . Equation (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) 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 the next period's mode and forward-looking variables, and the predetermined variables and instruments in period t . The matrix \mathbf{A}_{22j} is nonsingular for each $j \in N_j$.

The mode j_t follows a Markov process with the transition matrix $\mathbf{P} \equiv [P_{jk}]$.⁸ The shocks ε_t have mean zero and are i.i.d. with probability density φ , and we assume, without loss of generality, that ε_t is independent of j_t .⁹ We also assume that $\mathbf{C}_{1j}\varepsilon_t$ and $\mathbf{C}_{2k}\varepsilon_t$ 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 $\mathbf{p}_t = (p_{1t}, \dots, p_{n_t})'$ denote the true probability distribution of j_t in period t . We let $\mathbf{p}_{t+\tau|t}$ 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

$$\mathbf{p}_{t+1|t} = \mathbf{P}'\mathbf{p}_{t|t}. \tag{3}$$

We let the operator $E_t[\cdot]$ in the expression $E_t \mathbf{H}_{j_{t+1}} \mathbf{x}_{t+1}$ on the left-hand side of equation (2) denote expectations in period t , conditional on the policymaker's and the private sector's information in the beginning of period t , including $\mathbf{X}_t, \mathbf{i}_t$, and $\mathbf{p}_{t|t}$, but excluding j_t and ε_t . The maintained assumption is thus 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 . While we focus on the determination of the optimal policy instrument \mathbf{i}_t , our results also show how private sector choices as embodied in \mathbf{x}_t are affected by uncertainty and learning. The precise informational assumptions and the determination of $\mathbf{p}_{t|t}$ are specified below.

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

9. We can still incorporate additive mode-dependent shocks since the models allow mode-dependent intercepts (as well as mode-dependent standard deviations).

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

$$E_t \sum_{\tau=0}^{\infty} \delta^\tau L(\mathbf{X}_{t+\tau}, \mathbf{x}_{t+\tau}, \mathbf{i}_{t+\tau}, j_{t+\tau}), \quad (4)$$

where δ is a discount factor satisfying $0 < \delta < 1$, and the period loss, $L(\mathbf{X}_t, \mathbf{x}_t, \mathbf{i}_t, j_t)$, satisfies

$$L(\mathbf{X}_t, \mathbf{x}_t, \mathbf{i}_t, j_t) \equiv \begin{bmatrix} \mathbf{X}_t \\ \mathbf{x}_t \\ \mathbf{i}_t \end{bmatrix}' \mathbf{W}_{j_t} \begin{bmatrix} \mathbf{X}_t \\ \mathbf{x}_t \\ \mathbf{i}_t \end{bmatrix}, \quad (5)$$

where the matrix \mathbf{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 then add the term

$$\Xi_{t-1} \frac{1}{\delta} E_t \mathbf{H}_{j_t} \mathbf{x}_t \quad (6)$$

to the intertemporal loss function in period t . As we show below, the n_x vector Ξ_{t-1} is the vector of Lagrange multipliers for equation (2) from the optimization problem in period $t-1$. For the special case in which there are no forward-looking variables ($n_x = 0$), the model consists of equation (1) only, without the term $\mathbf{A}_{12j_{t+1}} \mathbf{x}_t$ the period loss function depends on \mathbf{X}_t , \mathbf{i}_t , and j_t only; and there is no role for the Lagrange multipliers Ξ_{t-1} or the term in equation (6).

1.2 Approximate MJLQ Models

While in this paper we start with an MJLQ model, the 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 (2007b) 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,

since an increasing time is spent in the vicinity of the steady state. Our MJLQ approximations are asymptotically valid for small shocks and persistent modes, since an increasing time is spent in the vicinity of each mode-dependent steady state. Thus, for slowly varying Markov chains, our MJLQ models provide accurate approximations of nonlinear models with Markov switching.

1.3 Types of Optimal Policies

We distinguish three cases: optimal policy when there is no learning (NL), adaptive optimal policy (AOP), and Bayesian optimal policy (BOP). By NL, we refer to a situation in which the policymaker and the aggregate private sector have a probability distribution $\mathbf{p}_{t|t}$ over the modes in period t and update the probability distribution in future periods using the transition matrix only, so the updating equation is

$$\mathbf{p}_{t+1|t+1} = \mathbf{P}'\mathbf{p}_{t|t} \tag{7}$$

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 $\mathbf{p}_{t|t}$ and equation (7). This is a variant of a case examined in Svensson and Williams (2007b).

By AOP, we refer to a situation in which 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 they 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 in which the policymaker acknowledges that the current instruments will affect future inference and updating of the probability distribution and takes this separate channel into account when calculating optimal policy. BOP thus includes optimal experimentation, whereby the policymaker may, for instance, 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.

1.3.1 Optimal policy with no learning

We first consider the NL case. Svensson and Williams (2007b) derive the equilibrium under commitment in a timeless perspective for the case in which \mathbf{X}_t , \mathbf{x}_t , and \mathbf{i}_t are observable in period t , j_t is unobservable, and the updating equation for $p_{t|t}$ is given by equation (7). Observations of \mathbf{X}_t , \mathbf{x}_t , and \mathbf{i}_t are then not used to update $p_{t|t}$.

It is useful to replace equation (2) by the two equivalent equations,

$$E_t \mathbf{H}_{j_{t+1}} \mathbf{x}_{t+1} = \mathbf{z}_t \quad (8)$$

and

$$0 = \mathbf{A}_{21j_t} \mathbf{X}_t + \mathbf{A}_{22j_t} \mathbf{x}_t - \mathbf{z}_t + \mathbf{B}_{2j_t} \mathbf{i}_t + \mathbf{C}_{2j_t} \varepsilon_t, \quad (9)$$

where we introduce the n_x vector of additional forward-looking variables, \mathbf{z}_t . Introducing this vector is a practical way of keeping track of the expectations term on the left-hand side of equation (2). Furthermore, it is practical to use equation (9) to solve \mathbf{x}_t as a function of \mathbf{X}_t , \mathbf{z}_t , \mathbf{i}_t , j_t , and ε_t :

$$\mathbf{x}_t = \tilde{\mathbf{x}}(\mathbf{X}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \varepsilon_t) \equiv \mathbf{A}_{22j_t}^{-1} \left(\mathbf{z}_t - \mathbf{A}_{21j_t} \mathbf{X}_t - \mathbf{B}_{2j_t} \mathbf{i}_t - \mathbf{C}_{2j_t} \varepsilon_t \right). \quad (10)$$

For a given j_t , this function is linear in \mathbf{X}_t , \mathbf{z}_t , \mathbf{i}_t , and ε_t .

To solve for the optimal decisions, we use the recursive saddlepoint method.¹⁰ We thus 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

$$E_t \tilde{L} \left(\tilde{\mathbf{X}}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t \right) \equiv \sum_j p_{j|t} \int \tilde{L} \left(\tilde{\mathbf{X}}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j, \varepsilon_t \right) \varphi(\varepsilon_t) d\varepsilon_t,$$

where $\tilde{\mathbf{X}}_t \equiv (\mathbf{X}'_t, \boldsymbol{\Xi}'_{t-1})'$ is the $(n_X + n_x)$ vector of extended predetermined variables (that is, including the n_x vector, $\boldsymbol{\Xi}_{t-1}$), γ_t is an n_x vector of

10. See Marcat and Marimon (1998), Svensson and Williams (2007b), and Svensson (2007) for details of the recursive saddlepoint method.

Lagrange multipliers, and $\varphi(\cdot)$ denotes a generic probability density function (for ε_t , the standard normal density function), and where

$$\begin{aligned} \tilde{L}\left(\tilde{\mathbf{X}}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t\right) &\equiv L\left(\mathbf{X}_t, \tilde{\mathbf{x}}\left(\mathbf{X}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \varepsilon_t\right), \mathbf{i}_t, j_t\right) \\ &- \gamma_t' \mathbf{z}_t + \Xi_{t-1}' \frac{1}{\delta} \mathbf{H}_{j_t} \tilde{\mathbf{x}}\left(\mathbf{X}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \varepsilon_t\right). \end{aligned} \quad (11)$$

As discussed in Svensson and Williams (2007b), the failure of the law of iterated expectations leads us to introduce a collection of value functions, $\hat{V}(\mathbf{s}_t, j)$, which condition on the mode, while the value function $\hat{V}(\mathbf{s}_t)$ averages over these and represents the solution of the dual optimization problem. The somewhat unusual Bellman equation for the dual problem can be written

$$\begin{aligned} \tilde{V}(\mathbf{s}_t) &\equiv E_t \hat{V}(\mathbf{s}_t, j_t) \equiv \sum_j p_{j|t} \hat{V}(\mathbf{s}_t, j) \\ &= \max_{\gamma_t} \min_{(\mathbf{z}_t, \mathbf{i}_t)} E_t \left[\tilde{L}(\tilde{\mathbf{X}}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t) + \delta \hat{V}(g(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}), j_{t+1}) \right] \\ &\equiv \max_{\gamma_t} \min_{(\mathbf{z}_t, \mathbf{i}_t)} \sum_j p_{j|t} \int \left[\tilde{L}(\tilde{\mathbf{X}}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j, \varepsilon_t) \right. \\ &\quad \left. + \delta \sum_k P_{jk} \hat{V}(g(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j, \varepsilon_t, k, \varepsilon_{t+1}), k) \right] \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\boldsymbol{\varepsilon}_t d\boldsymbol{\varepsilon}_{t+1}. \end{aligned} \quad (12)$$

where $\mathbf{s}_t \equiv (\tilde{\mathbf{X}}_t', \mathbf{p}'_{t|t})'$ denotes the perceived state of the economy (it includes the perceived probability distribution, $\mathbf{p}_{t|t}$, but not the true mode) and (\mathbf{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 \mathbf{s}_t . Also, in the Bellman equation we require that all the choice variables respect the information constraints, and they thus depend on the perceived state \mathbf{s}_t but not directly on the mode j .

The optimization is subject to the transition equation for \mathbf{X}_t ,

$$\mathbf{X}_{t+1} = \mathbf{A}_{11j_{t+1}} \mathbf{X}_t + \mathbf{A}_{12j_{t+1}} \tilde{\mathbf{x}}(\mathbf{X}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \varepsilon_t) + \mathbf{B}_{1j_{t+1}} \mathbf{i}_t + \mathbf{C}_{1j_{t+1}} \varepsilon_{t+1}, \quad (13)$$

where we have substituted $\tilde{\mathbf{x}}(\mathbf{X}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \varepsilon_t)$ for \mathbf{x}_t ; the new dual transition equation for Ξ_t ,

$$\Xi_t = \gamma_t, \quad (14)$$

and the transition equation (7) for $\mathbf{p}_{t|t}$. Combining equations, we have the transition for \mathbf{s}_t ,

$$\mathbf{s}_{t+1} \equiv \begin{bmatrix} \mathbf{X}_{t+1} \\ \Xi_t \\ \mathbf{p}_{t+1|t+1} \end{bmatrix} = g(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \quad (15)$$

$$\equiv \begin{bmatrix} \mathbf{A}_{11j_{t+1}} \mathbf{X}_t + \mathbf{A}_{12j_{t+1}} \tilde{\mathbf{x}}(\mathbf{X}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \varepsilon_t) + \mathbf{B}_{1j_{t+1}} \mathbf{i}_t + \mathbf{C}_{1j_{t+1}} \varepsilon_{t+1} \\ \gamma_t \\ \mathbf{P}' \mathbf{p}_{t|t} \end{bmatrix}.$$

It is straightforward to see that the solution of the dual optimization problem (equation 12) is linear in $\tilde{\mathbf{X}}_t$ for given $\mathbf{p}_{t|t}$, j_t ,

$$\begin{bmatrix} \mathbf{z}_t \\ \mathbf{i}_t \\ \gamma_t \end{bmatrix} = \begin{bmatrix} \mathbf{z}(\mathbf{s}_t) \\ \mathbf{i}(\mathbf{s}_t) \\ \gamma(\mathbf{s}_t) \end{bmatrix} = F(\mathbf{p}_{t|t}) \tilde{\mathbf{X}}_t \equiv \begin{bmatrix} F_z(\mathbf{p}_{t|t}) \\ F_i(\mathbf{p}_{t|t}) \\ F_\gamma(\mathbf{p}_{t|t}) \end{bmatrix} \tilde{\mathbf{X}}_t, \quad (16)$$

$$\mathbf{x}_t = \mathbf{x}(\mathbf{s}_t, j_t, \varepsilon_t) \equiv \tilde{\mathbf{x}}(\mathbf{X}_t, \mathbf{z}(\mathbf{s}_t), \mathbf{i}(\mathbf{s}_t), j_t, \varepsilon_t) \quad (17)$$

$$\equiv F_{\tilde{\mathbf{x}}\tilde{\mathbf{X}}}(\mathbf{p}_{t|t}, j_t) \tilde{\mathbf{X}}_t + F_{\mathbf{x}\varepsilon}(\mathbf{p}_{t|t}, j_t) \varepsilon_t.$$

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

$$\mathbf{s}_{t+1} = \hat{g}(\mathbf{s}_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \equiv g(\mathbf{s}_t, \mathbf{z}(\mathbf{s}_t), \mathbf{i}(\mathbf{s}_t), \gamma(\mathbf{s}_t), j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}). \quad (18)$$

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

$$\tilde{V}(\mathbf{s}_t) \equiv \tilde{\mathbf{X}}_t' \tilde{V}_{\tilde{\mathbf{X}}\tilde{\mathbf{X}}}(\mathbf{p}_{t|t}) \tilde{\mathbf{X}}_t + w(\mathbf{p}_{t|t}).$$

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

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

The function $\hat{V}(\mathbf{s}_t, j)$ is also quadratic in $\tilde{\mathbf{X}}_t$ for given $\mathbf{p}_{t|t}$ and j ,

$$\hat{V}(\mathbf{s}_t, j) \equiv \tilde{\mathbf{X}}_t' \hat{V}_{\mathbf{XX}}(\mathbf{p}_{t|t}, j) \tilde{\mathbf{X}}_t + \hat{w}(\mathbf{p}_{t|t}, j).$$

It follows that we have

$$\tilde{V}_{\mathbf{XX}}(\mathbf{p}_{t|t}) \equiv \sum_j p_{j|t} \hat{V}_{\mathbf{XX}}(\mathbf{p}_{t|t}, j);$$

$$w(\mathbf{p}_{t|t}) \equiv \sum_j p_{j|t} \hat{w}(\mathbf{p}_{t|t}, j).$$

Although we find the optimal policies from the dual problem, we use the value function for the primal problem (with the original, unmodified loss function) to measure true expected losses. This value function, with the period loss function $E_t L(\mathbf{X}_t, \mathbf{x}_t, \mathbf{i}_t, j_t)$ rather than $E_t \tilde{L}(\tilde{\mathbf{X}}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t)$, satisfies

$$\begin{aligned} V(\mathbf{s}_t) &\equiv \tilde{V}(\mathbf{s}_t) - \Xi'_{t-1} \frac{1}{\delta} \sum_j p_{j|t} \mathbf{H}_j \int \mathbf{x}(\mathbf{s}_t, j, \varepsilon_t) \varphi(\varepsilon_t) d\varepsilon_t \\ &= \tilde{V}(\mathbf{s}_t) - \Xi'_{t-1} \frac{1}{\delta} \sum_j p_{j|t} \mathbf{H}_j \mathbf{x}(\mathbf{s}_t, j, 0). \end{aligned} \tag{19}$$

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

$$V(\mathbf{s}_t) \equiv \tilde{\mathbf{X}}_t' V_{\mathbf{XX}}(\mathbf{p}_{t|t}) \tilde{\mathbf{X}}_t + w(\mathbf{p}_{t|t}),$$

where the scalar $w(\mathbf{p}_{t|t})$ in the primal value function is identical to that in the dual value function. This is the value function conditional on $\tilde{\mathbf{X}}_t$ and $\mathbf{p}_{t|t}$ after $\tilde{\mathbf{X}}_t$ has been observed but before \mathbf{x}_t has been observed, taking into account that j_t and ε_t are not observed. Hence, the second

term on the right-hand side of equation (19) contains the expectation of $\mathbf{H}_t \mathbf{x}_t$ conditional on that information.¹¹

In Svensson and Williams (2007a, 2007b), we 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

$$V(\mathbf{s}_t) \equiv \sum_j p_{j|t} \bar{V}(\mathbf{s}_t, j),$$

where the conditional value function, $\bar{V}(\mathbf{s}_t, j)$, satisfies

$$\bar{V}(\mathbf{s}_t, j) = \int \left[L(\mathbf{X}_t, \mathbf{x}(\mathbf{s}_t, j, \varepsilon_t), \mathbf{i}(\mathbf{s}_t), j) \right. \\ \left. + \delta \sum_k P_{jk} \bar{V}(\hat{\mathbf{g}}(\mathbf{s}_t, j, \varepsilon_t, k, \varepsilon_{t+1}), k) \right] \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1}, \quad (j \in N_j). \quad (20)$$

1.3.2 Adaptive optimal policy

Consider now the case of adaptive optimal policy, in which the policymaker uses the same policy function as in the no-learning case, but each period updates the probabilities on which this policy is conditioned. 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 $\mathbf{C}_{2j_t} \geq 0$ and that both ε_t and j_t are unobservable. The estimate $\mathbf{p}_{t|t}$ is the result of Bayesian updating, using all information available, but the optimal policy in period t is computed under the perceived updating equation (7). That is, we disregard the fact that the policy choice will affect future $\mathbf{p}_{t+\tau|t+\tau}$ and that future expected loss will change when $\mathbf{p}_{t+\tau|t+\tau}$ changes. Under the assumption that the expectations on the left-hand side of equation (2) are conditional on equation (7), the variables \mathbf{z}_t , \mathbf{i}_t , γ_t , and \mathbf{x}_t in period t are still determined by equations (16) and (17).

To determine the updating equation for $\mathbf{p}_{t|t}$, we specify an explicit sequence of information revelation as follows, in nine steps. The timing assumptions are necessary to spell out the appropriate conditioning for decisions and updating of beliefs.

11. To be precise, the observation of \mathbf{X}_t , which depends on $\mathbf{C}_{1j_t \varepsilon_t}$, allows some inference of ε_t , $\varepsilon_{t|t}$. The variable \mathbf{x}_t will depend on j_t and on ε_t , but on ε_t only through $\mathbf{C}_{2j_t \varepsilon_t}$. By assumption, $\mathbf{C}_{1j_t \varepsilon_t}$ and $\mathbf{C}_{2k \varepsilon_t}$ are independent. Hence, any observation of \mathbf{X}_t and $\mathbf{C}_{1j_t \varepsilon_t}$ does not convey any information about $\mathbf{C}_{2j_t \varepsilon_t}$, so $E_t \mathbf{C}_{2j_t \varepsilon_t} = 0$.

First, the policymaker and the private sector enter period t with the prior $\mathbf{p}_{t|t-1}$. They know \mathbf{X}_{t-1} , $\mathbf{x}_{t-1} = \mathbf{x}(\mathbf{s}_{t-1}, j_{t-1}, \boldsymbol{\varepsilon}_{t-1})$, $\mathbf{z}_{t-1} = \mathbf{z}(\mathbf{s}_{t-1})$, $\mathbf{i}_{t-1} = \mathbf{i}(\mathbf{s}_{t-1})$, and $\boldsymbol{\Xi}_{t-1} = \boldsymbol{\gamma}(\mathbf{s}_{t-1})$ from the previous period.

Second, the mode j_t and the vector of shocks $\boldsymbol{\varepsilon}_t$ are realized in the beginning of period t . The vector of predetermined variables \mathbf{X}_t is then realized according to equation (1).

Third, the policymaker and the private sector observe \mathbf{X}_t . They then know that $\tilde{\mathbf{X}}_t \equiv (\mathbf{X}'_t, \boldsymbol{\Xi}'_{t-1})'$. They do not observe j_t or $\boldsymbol{\varepsilon}_t$.

Fourth, the policymaker and the private sector update the prior $\mathbf{p}_{t|t-1}$ to the posterior $\mathbf{p}_{t|t}$ according to Bayes' theorem and the updating equation

$$p_{j|t} = \frac{\varphi(\mathbf{X}_t | j_t = j, \mathbf{X}_{t-1}, \mathbf{x}_{t-1}, \mathbf{i}_{t-1}, \mathbf{p}_{t|t-1})}{\varphi(\mathbf{X}_t | \mathbf{X}_{t-1}, \mathbf{x}_{t-1}, \mathbf{i}_{t-1}, \mathbf{p}_{t|t-1})} p_{j|t-1}, \quad (j \in N_j), \quad (21)$$

where again $\varphi(\cdot)$ denotes a generic density function.¹² Then the policymaker and the private sector know that $\mathbf{s}_t \equiv (\tilde{\mathbf{X}}'_t, \mathbf{p}'_{t|t})'$.

Fifth, the policymaker solves the dual optimization problem, determines $\mathbf{i}_t = \mathbf{i}(\mathbf{s}_t)$, and implements or announces the instrument setting \mathbf{i}_t .

Sixth, the private sector and policymaker form their expectations,

$$\mathbf{z}_t = E_t \mathbf{H}_{j_{t+1}} \mathbf{x}_{t+1} \equiv E \left[\mathbf{H}_{j_{t+1}} \mathbf{x}_{t+1} | \mathbf{s}_t \right].$$

In equilibrium, these expectations will be determined by equation (16). These expectations are by assumption formed before \mathbf{x}_t is observed. The private sector and the policymaker know that \mathbf{x}_t will, in equilibrium, be determined in the next step according to equation (17). Hence, they can form expectations of the soon-to-be determined \mathbf{x}_t conditional on $j_t = j$,¹³

$$\mathbf{x}_{j|t} = \mathbf{x}(\mathbf{s}_t, j, 0). \quad (22)$$

12. The policymaker and private sector can also estimate the shocks $\boldsymbol{\varepsilon}_{t|t}$ as $\boldsymbol{\varepsilon}_{t|t} = \sum_j p_{j|t} \boldsymbol{\varepsilon}_{j|t}$, where $\boldsymbol{\varepsilon}_{j|t} \equiv \mathbf{X}_t - \mathbf{A}_{11j} \mathbf{X}_{t-1} - \mathbf{A}_{12j} \mathbf{x}_{t-1} - \mathbf{B}_{1j} \mathbf{i}_{t-1}$ ($j \in N_j$). However, because of the assumed independence of $\mathbf{C}_{1j} \boldsymbol{\varepsilon}_t$ and $\mathbf{C}_{2k} \boldsymbol{\varepsilon}_t$, $j, k \in N_j$, we do not need to keep track of $\boldsymbol{\varepsilon}_{j|t}$.

13. Note that 0 instead of $\boldsymbol{\varepsilon}_{j|t}$ enters above. The inference $\boldsymbol{\varepsilon}_{j|t}$ above is inference about $\mathbf{C}_{1j} \boldsymbol{\varepsilon}_t$, whereas \mathbf{x}_t depends on $\boldsymbol{\varepsilon}_t$ through $\mathbf{C}_{2j} \boldsymbol{\varepsilon}_t$. Since we assume that $\mathbf{C}_{1j} \boldsymbol{\varepsilon}_t$ and $\mathbf{C}_{2j} \boldsymbol{\varepsilon}_t$ are independent, there is no inference of $\mathbf{C}_{2j} \boldsymbol{\varepsilon}_t$ from observing \mathbf{X}_t . Hence, $E_t \mathbf{C}_{2j} \boldsymbol{\varepsilon}_t \equiv 0$. Because of the linearity of \mathbf{x}_t in $\boldsymbol{\varepsilon}_t$, the integration of \mathbf{x}_t over $\boldsymbol{\varepsilon}_t$ results in $\mathbf{x}(\mathbf{s}_t, j, 0)$.

The private sector and the policymaker can also infer Ξ_t from

$$\Xi_t = \gamma(\mathbf{s}_t). \quad (23)$$

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

$$\mathbf{z}_t = \mathbf{z}(\mathbf{s}_t) = E_t \left[\mathbf{H}_{j_{t+1}} \mathbf{x}_{t+1} | \mathbf{s}_t \right] = \sum_{j,k} P_{jk} p_{j|t} \mathbf{H}_k \mathbf{x}_{k,t+1|jt}, \quad (24)$$

where

$$\begin{aligned} \mathbf{x}_{k,t+1|jt} &= \int \mathbf{x} \left(\left[\begin{array}{c} \mathbf{A}_{11k} \mathbf{X}_t + \mathbf{A}_{12k} \mathbf{x}(\mathbf{s}_t, j, \epsilon_t) + \mathbf{B}_{1k} \mathbf{i}(\mathbf{s}_t) \\ \Xi_t \\ \mathbf{P}' \mathbf{p}_{t|t} \end{array} \right], k, \epsilon_{t+1} \right) \varphi(\epsilon_t) \varphi(\epsilon_{t+1}) d\epsilon_t d\epsilon_{t+1} \\ &= \mathbf{x} \left(\left[\begin{array}{c} \mathbf{A}_{11k} \mathbf{X}_t + \mathbf{A}_{12k} \mathbf{x}(\mathbf{s}_t, j, 0) + \mathbf{B}_{1k} \mathbf{i}(\mathbf{s}_t) \\ \Xi_t \\ \mathbf{P}' \mathbf{p}_{t|t} \end{array} \right], k, \mathbf{0} \right), \end{aligned}$$

and where we have exploited the linearity of $\mathbf{x}_t = \mathbf{x}(\mathbf{s}_t, j_t, \epsilon_t)$ and $\mathbf{x}_{t+1} = \mathbf{x}(\mathbf{s}_{t+1}, j_{t+1}, \epsilon_{t+1})$ in ϵ_t and ϵ_{t+1} . Under AOP, \mathbf{z}_t is formed conditional on the belief that the probability distribution in period $t+1$ will be given by $\mathbf{p}_{t+1|t+1} = \mathbf{P}' \mathbf{p}_{t|t}$, not by the true updating equation that we are about to specify.

Seventh, after the expectations \mathbf{z}_t have been formed, \mathbf{x}_t is determined as a function of \mathbf{X}_t , \mathbf{z}_t , \mathbf{i}_t , j_t , and ϵ_t by equation (10).

Eighth, the policymaker and the private sector then use the observed \mathbf{x}_t to update $\mathbf{p}_{t|t}$ to the new posterior $\mathbf{p}_{t|t}^+$ according to Bayes' theorem, via the updating equation

$$\mathbf{p}_{j|t}^+ = \frac{\varphi(\mathbf{x}_t | j_t = j, \mathbf{X}_t, \mathbf{z}_t, \mathbf{i}_t, \mathbf{p}_{t|t})}{\varphi(\mathbf{x}_t | \mathbf{X}_t, \mathbf{z}_t, \mathbf{i}_t, \mathbf{p}_{t|t})} p_{j|t}, \quad (j \in N_j). \quad (25)$$

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

$$\mathbf{p}_{t+1|t} = \mathbf{P}' \mathbf{p}_{t|t}^+ \tag{26}$$

In the beginning of period $t + 1$, the mode j_{t+1} and the vector of shocks $\boldsymbol{\varepsilon}_{t+1}$ are realized, and \mathbf{X}_{t+1} is determined by equation (1) and observed by the policymaker and the private sector. The sequence of the nine steps above then repeats itself. For more detail on the explicit densities in the updating equations (21) and (25), see Svensson and Williams (2007a).

The transition equation for $\mathbf{p}_{t+1|t+1}$ can be written

$$\mathbf{p}_{t+1|t+1} = Q(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \boldsymbol{\varepsilon}_t, j_{t+1}, \boldsymbol{\varepsilon}_{t+1}), \tag{27}$$

where $Q(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \boldsymbol{\varepsilon}_t, j_{t+1}, \boldsymbol{\varepsilon}_{t+1})$ is defined by the combination of equation (21) for period $t + 1$ with equations (13) and (26). The equilibrium transition equation for the full state vector is then given by

$$\begin{aligned} \mathbf{s}_{t+1} &\equiv \begin{bmatrix} \mathbf{X}_{t+1} \\ \boldsymbol{\Xi}_t \\ \mathbf{p}_{t+1|t+1} \end{bmatrix} = \bar{g}(\mathbf{s}_t, j_t, \boldsymbol{\varepsilon}_t, j_{t+1}, \boldsymbol{\varepsilon}_{t+1}) \\ &\equiv \begin{bmatrix} \mathbf{A}_{11j_{t+1}} \mathbf{X}_t + \mathbf{A}_{12j_{t+1}} \mathbf{x}(\mathbf{s}_t, j_t, \boldsymbol{\varepsilon}_t) + \mathbf{B}_{1j_{t+1}} \mathbf{i}(\mathbf{s}_t) + \mathbf{C}_{1j_{t+1}} \boldsymbol{\varepsilon}_{t+1} \\ \gamma(\mathbf{s}_t) \\ Q(\mathbf{s}_t, \mathbf{z}(\mathbf{s}_t), \mathbf{i}(\mathbf{s}_t), j_t, \boldsymbol{\varepsilon}_t, j_{t+1}, \boldsymbol{\varepsilon}_{t+1}) \end{bmatrix}, \end{aligned} \tag{28}$$

where the third row is given by the true updating equation (27) together with the policy function (16). Thus, in this AOP case, there is a distinction between the “perceived” transition equation (15) and the equilibrium transition equation (18), both of which include the perceived updating equation (7) in the bottom block, and the “true” equilibrium transition equation (28), which replaces the perceived updating equation (7) with the true updating equation (27).

Note that $V(\mathbf{s}_t)$ in equation (19), which is subject to the perceived transition equation (15), does not give the true (unconditional) value function for the AOP case. This is instead given by

$$\bar{V}(\mathbf{s}_t) \equiv \sum_j p_{j|t} \bar{V}(\mathbf{s}_t, j),$$

where the true conditional value function, $\bar{V}(\mathbf{s}_t, j_t)$, satisfies

$$\bar{V}(\mathbf{s}_t, j) = \int \left[L(\mathbf{X}_t, \mathbf{x}(\mathbf{s}_t, j, \mathbf{e}_t), \mathbf{i}(\mathbf{s}_t), j) \right. \\ \left. + \delta \sum_k P_{jk} \bar{V}(\bar{\mathbf{g}}(\mathbf{s}_t, j, \mathbf{e}_t, k, \mathbf{e}_{t+1}), k) \right] \varphi(\mathbf{e}_t) \varphi(\mathbf{e}_{t+1}) d\mathbf{e}_t d\mathbf{e}_{t+1}, \quad (j \in N_j). \quad (29)$$

That is, the true value function $\bar{V}(\mathbf{s}_t)$ takes into account the true updating equation for $\mathbf{p}_{t|t}$, equation (27), whereas the optimal policy, equation (16), and the perceived value function, $V(\mathbf{s}_t)$ in equation (19), are conditional on the perceived updating equation (7) and thereby the perceived transition equation (15). Also, $\bar{V}(\mathbf{s}_t)$ is the value function after $\bar{\mathbf{X}}_t$ has been observed but before \mathbf{x}_t is observed, so it is conditional on $\mathbf{p}_{t|t}$ rather than on $\mathbf{p}_{t|t}^+$. Since the full transition equation (28) is no longer linear given the belief updating in equation (27), the true value function $\bar{V}(\mathbf{s}_t)$ is no longer quadratic in $\bar{\mathbf{X}}_t$ for given $\mathbf{p}_{t|t}$. 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 (2007a), the difference between the true updating equation for $\mathbf{p}_{t+1|t+1}$, (27), and the perceived updating equation (7) is that in the true updating equation, $\mathbf{p}_{t+1|t+1}$ becomes a random variable from the point of view of period t , with mean equal to $\mathbf{p}_{t+1|t}$. This is because $\mathbf{p}_{t+1|t+1}$ depends on the realization of j_{t+1} and \mathbf{e}_{t+1} . Bayesian updating thus induces a mean-preserving spread over beliefs, which in turn sheds light on the gains from learning. If the conditional value function $\bar{V}(\mathbf{s}_t, j_t)$ under NL is concave in $\mathbf{p}_{t|t}$ for given $\bar{\mathbf{X}}_t$ 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 is likely that \bar{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 $\mathbf{p}_{t|t}$ that is, to achieve a mean-preserving spread that might further reduce the expected future loss.¹⁴ This amounts to optimal experimentation.

14. Kiefer (1989) examines the properties of a value function, including concavity, under Bayesian learning for a simpler model without forward-looking variables.

1.3.3 Bayesian optimal policy

Finally, we consider the BOP case, in which optimal policy is determined while taking the updating equation (27) into account. That is, we now allow the policymaker to choose \mathbf{i}_t taking into account that his actions will affect $\mathbf{p}_{t+1|t+1}$, which in turn will affect future expected losses. In particular, experimentation is allowed and is optimally chosen. Hence, for the BOP case, there is no distinction between the “perceived” and “true” transition equations.

The transition equation for the BOP case is

$$\mathbf{s}_{t+1} \equiv \begin{bmatrix} \mathbf{X}_{t+1} \\ \Xi_t \\ \mathbf{p}_{t+1|t+1} \end{bmatrix} = g(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \quad (30)$$

$$\equiv \begin{bmatrix} \mathbf{A}_{11j_{t+1}} \mathbf{X}_t + \mathbf{A}_{12j_{t+1}} \tilde{\mathbf{x}}(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \varepsilon_t) + \mathbf{B}_{1j_{t+1}} \mathbf{i}_t + \mathbf{C}_{1j_{t+1}} \varepsilon_{t+1} \\ \gamma_t \\ Q(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \end{bmatrix}.$$

Then the dual optimization problem can be written as equation (12) subject to the above transition equation (30). Matters simplify somewhat in the Bayesian case, however, as we do not need to compute the conditional value functions $\hat{V}(\mathbf{s}_t, j_t)$, which were required in the AOP case given the failure of the law of iterated expectations. The second term on the right-hand side of equation (12) can be written as

$$E_t \hat{V}(\mathbf{s}_{t+1}, j_{t+1}) \equiv E \left[\hat{V}(\mathbf{s}_{t+1}, j_{t+1}) \middle| \mathbf{s}_t \right].$$

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

$$E \left[\hat{V}(\mathbf{s}_{t+1}, j_{t+1}) \middle| \mathbf{s}_t \right] = E \left[\tilde{V}(\mathbf{s}_{t+1}) \middle| \mathbf{s}_t \right].$$

See Svensson and Williams (2007a) for a proof.

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

$$\begin{aligned}
\tilde{V}(\mathbf{s}_t) &= \max_{\gamma_t} \min_{(\mathbf{z}_t, \mathbf{i}_t)} E_t \left[\tilde{L}(\tilde{\mathbf{X}}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t) + \delta \tilde{V}(g(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1})) \right] \\
&\equiv \max_{\gamma_t} \min_{(\mathbf{z}_t, \mathbf{i}_t)} \sum_j P_{j|t} \int \left[\tilde{L}(\tilde{\mathbf{X}}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j, \varepsilon_t) \right. \\
&\quad \left. + \delta \sum_k P_{jk} \tilde{V}(g(\mathbf{s}_t, \mathbf{z}_t, \mathbf{i}_t, \gamma_t, j, \varepsilon_t, k, \varepsilon_{t+1})) \right] \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1}, \tag{31}
\end{aligned}$$

where the transition equation is given by equation (30).

The solution to the optimization problem can be written

$$\tilde{\mathbf{i}}_t \equiv \begin{bmatrix} \mathbf{z}_t \\ \mathbf{i}_t \\ \gamma_t \end{bmatrix} = \tilde{\mathbf{i}}(\mathbf{s}_t) \equiv \begin{bmatrix} \mathbf{z}(\mathbf{s}_t) \\ \mathbf{i}(\mathbf{s}_t) \\ \gamma(\mathbf{s}_t) \end{bmatrix} = F(\tilde{\mathbf{X}}_t, \mathbf{p}_{t|t}) \equiv \begin{bmatrix} F_z(\tilde{\mathbf{X}}_t, \mathbf{p}_{t|t}) \\ F_i(\tilde{\mathbf{X}}_t, \mathbf{p}_{t|t}) \\ F_\gamma(\tilde{\mathbf{X}}_t, \mathbf{p}_{t|t}) \end{bmatrix}, \tag{32}$$

$$\mathbf{x}_t = \mathbf{x}(\mathbf{s}_t, j_t, \varepsilon_t) \equiv \tilde{\mathbf{x}}(\tilde{\mathbf{X}}_t, \mathbf{z}(\mathbf{s}_t), \mathbf{i}(\mathbf{s}_t), j_t, \varepsilon_t) \equiv F_x(\tilde{\mathbf{X}}_t, \mathbf{p}_{t|t}, j_t, \varepsilon_t). \tag{33}$$

Because of the nonlinearity of equations (27) and (30), the solution is no longer linear in $\tilde{\mathbf{X}}_t$ for given $\mathbf{p}_{t|t}$. The dual value function, $\tilde{V}(\mathbf{s}_t)$, is no longer quadratic in $\tilde{\mathbf{X}}_t$ for given $\mathbf{p}_{t|t}$. The value function of the primal problem, $V(\mathbf{s}_t)$, is given by, equivalently, equation (19); equation (29) with the equilibrium transition equation (28) and with the solution (32); or

$$V(\mathbf{s}_t) = \sum_j P_{j|t} \int \left[L(\tilde{\mathbf{X}}_t, \mathbf{x}(\mathbf{s}_t, j, \varepsilon_t), \mathbf{i}(\mathbf{s}_t), j) \right. \\
\left. + \delta \sum_k P_{jk} V(\bar{g}(\mathbf{s}_t, j, \varepsilon_t, k, \varepsilon_{t+1})) \right] \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1}. \tag{34}$$

It is also no longer quadratic in $\tilde{\mathbf{X}}_t$ for given $\mathbf{p}_{t|t}$. More complex and detailed numerical methods are thus 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 here 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.

2. LEARNING AND EXPERIMENTATION IN A SIMPLE NEW-KEYNESIAN MODEL

We consider the benchmark standard New-Keynesian model, consisting of a New-Keynesian Phillips curve and a consumption Euler equation:¹⁵

$$\pi_t = \delta E_t \pi_{t+1} + \gamma_{j_t} y_t + c_\pi \varepsilon_{\pi t}; \quad (35)$$

$$y_t = E_t y_{t+1} - \sigma_{j_t} (i_t - E_t \pi_{t+1}) + c_y \varepsilon_{y_t} + c_g g_t; \quad (36)$$

$$g_{t+1} = \rho g_t + \varepsilon_{g,t+1}. \quad (37)$$

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 coming from variation in preferences, government spending, or the underlying efficient level of output. Woodford (2003) 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, some components of the shocks need to be unobservable in order for there to be a nontrivial inference problem for agents. We have 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) 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 nonhomothetic preferences coupled with shifts in output (in which case the preferences themselves would not shift, 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,

15. See Woodford (2003) for an exposition.

as we consider the usual approximations around a zero inflation rate and an efficient level of output.

2.1 Optimal Policy: No Learning, Adaptive Optimal Policy, and Bayesian Optimal Policy

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:

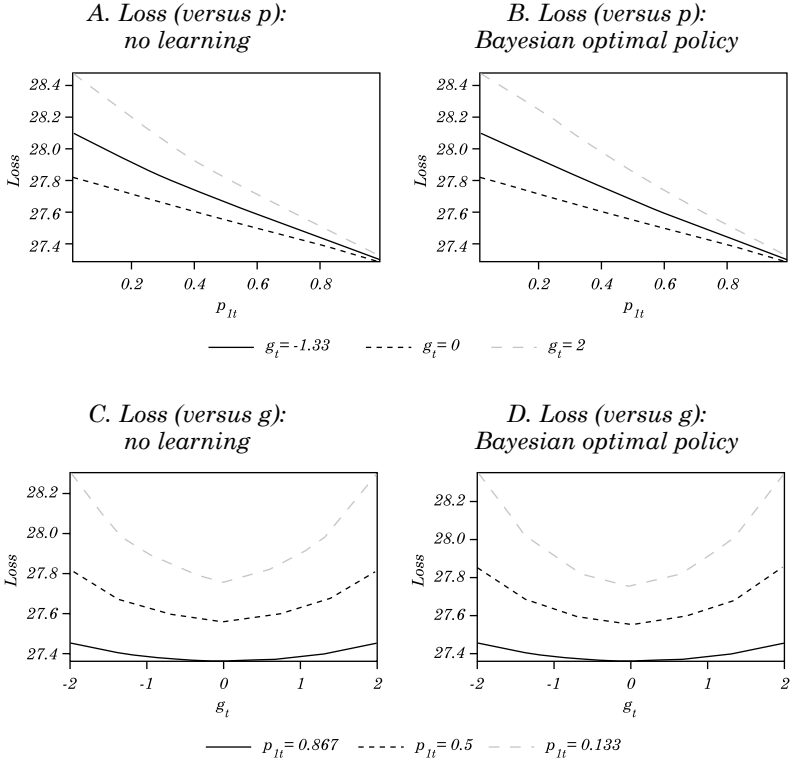
$$L_t = \pi_t^2 + \lambda_j y_t^2 + \mu i_t^2. \quad (38)$$

We set the following parameters, mostly following Woodford's (2003) calibration, as follows: $\gamma_1 = 0.024$, $\gamma_2 = 0.075$, $\sigma_1 = 1.000/0.157 = 6.370$, $\sigma_2 = 1.0$, $c_\pi = c_y = c_g = 0.5$, 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 (2003), while the two modes represent reasonable alternatives. Mode 1 is Woodford's benchmark case; 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

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

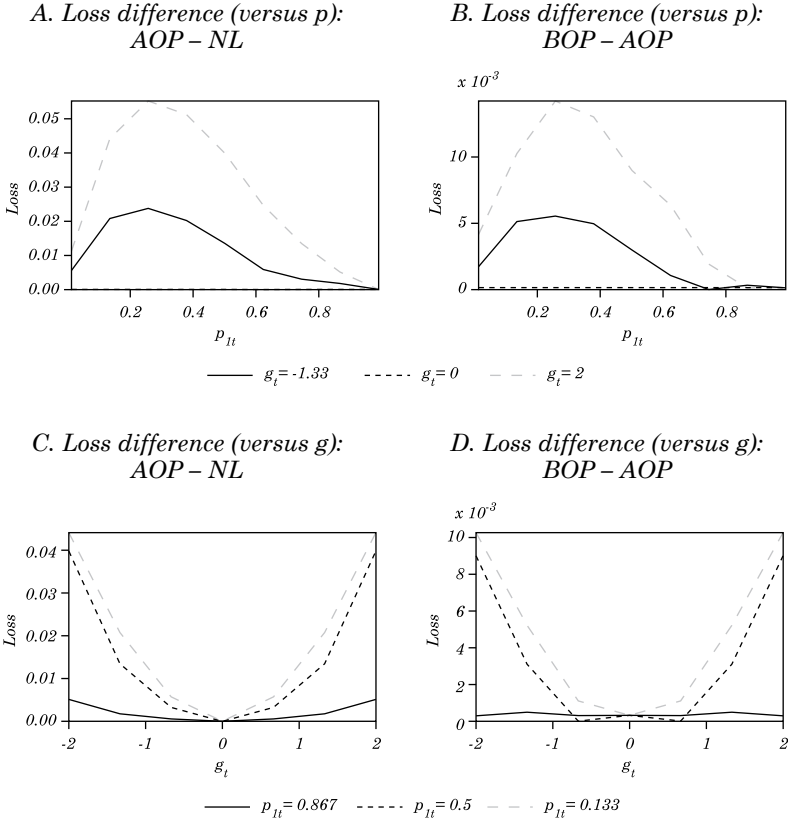
We have two forward-looking variables, $\mathbf{x}_t \equiv (\pi_t, y_t)'$, and consequently two Lagrange multipliers, $\Xi_{t-1} \equiv (\Xi_{\pi,t-1}, \Xi_{y,t-1})'$. We have one predetermined variable ($\mathbf{X}_t \equiv g_t$) and the estimated mode probabilities, $\mathbf{p}_{t|t} \equiv (p_{1t|t}, p_{2t|t})'$ (of which we only need keep track of one, $p_{1t|t}$). Thus, the value and policy functions, $V(\mathbf{s}_t)$ and $i(\mathbf{s}_t)$, are all four dimensional: $\mathbf{s}_t = (g_t, \Xi'_{t-1}, p_{1t|t})'$. We are therefore 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 1 shows losses under NL and BOP as functions of p_{1t} and g_t . Figure 2 shows the difference between losses under NL, AOP, and BOP. Figures 3 and 4 show the corresponding policy functions and their differences.

Figure 1. Losses from No Learning and Bayesian Optimal Policy



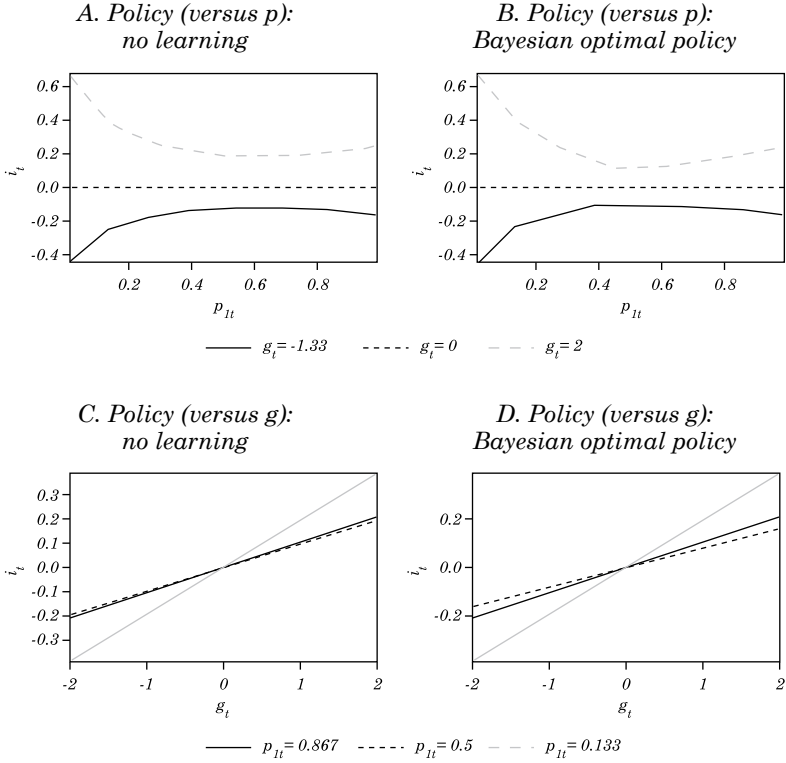
Source: Authors' calculations.

Figure 2. Differences in Losses from No Learning, Adaptive Optimal Policy, and Bayesian Optimal Policy




Source: Authors' calculations.

Figure 3. Optimal Policies under No Learning and Bayesian Optimal Policy



Source: Authors' calculations.

Figure 4. Differences in Policies under No Learning and Bayesian Optimal Policy



Source: Authors' calculations.

In Svensson and Williams (2007a) we show that learning implies a mean-preserving spread of the random variable $\mathbf{p}_{t+1|t+1}$ (which under learning is 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, figure 1 illustrates that the value function is actually slightly convex in p_{1t} , so learning is not beneficial here. In contrast, the value function is concave and learning is beneficial in a backward-looking example in Svensson and Williams (2007a).

Consequently, AOP gives higher losses than NL, as shown in figure 2. Furthermore, somewhat surprisingly, 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 (2007a).

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. Backward-looking models generally find 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.) Neither of these conclusions holds with forward-looking models. Under our assumption of symmetric information and beliefs between the private sector and the policymaker, both the private sector and the policymaker learn. 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, as it may induce further volatility in agents' beliefs.¹⁶

3. LEARNING IN AN ESTIMATED EMPIRICAL NEW-KEYNESIAN MODEL

The previous section focused on a simple small model to explore the impacts of learning and experimentation. Since computing

16. In the forward-looking case, we solve saddlepoint problems, and moving from AOP to BOP expands the feasible set for both the minimizing and maximizing choices.

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, using a version of Lindé's (2005) model that we estimated in Svensson and Williams (2007b). 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. We therefore focus on the impact of learning on policy and compare NL and AOP. In Svensson and Williams (2007b), we computed the optimal policy under no learning, and here we see how inference on the mode affects the dynamics of output, inflation, and interest rates.

3.1 The Model

The structural model is a mode-dependent simplification of Lindé's (2005) model of the U.S. economy and is given by

$$\begin{aligned}\pi_t &= \omega_{ff} E_t \pi_{t+1} + (1 - \omega_{ff}) \pi_{t-1} + \gamma_j y_t + c_{\pi j} \varepsilon_{\pi t}; \\ y_t &= \beta_{ff} E_t y_{t+1} + (1 - \beta_{ff}) [\beta_{yj} y_{t-1} + (1 - \beta_{yj}) y_{t-2}] \\ &\quad - \beta_{rj} (i_t - E_t \pi_{t+1}) + c_{yj} \varepsilon_{yt}.\end{aligned}\tag{39}$$

Here $j \in \{1, 2\}$ indexes the mode, and the shocks, $\varepsilon_{\pi t}$, ε_{yt} , and ε_{it} , are independent standard normal random variables. In particular, we consider a two-mode MJLQ model in which one mode has forward- and backward-looking elements and 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 we do not include that here since our focus is on optimal policy.

In Svensson and Williams (2007b), we estimate the model on U.S. data using Bayesian methods. The maximum posterior estimates are given in table 1, with the unconditional expectation of the coefficients for comparison. Apart from the forward-looking terms (which 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 \mathbf{P} and its implied stationary distribution $\bar{\mathbf{p}}$ are given by

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

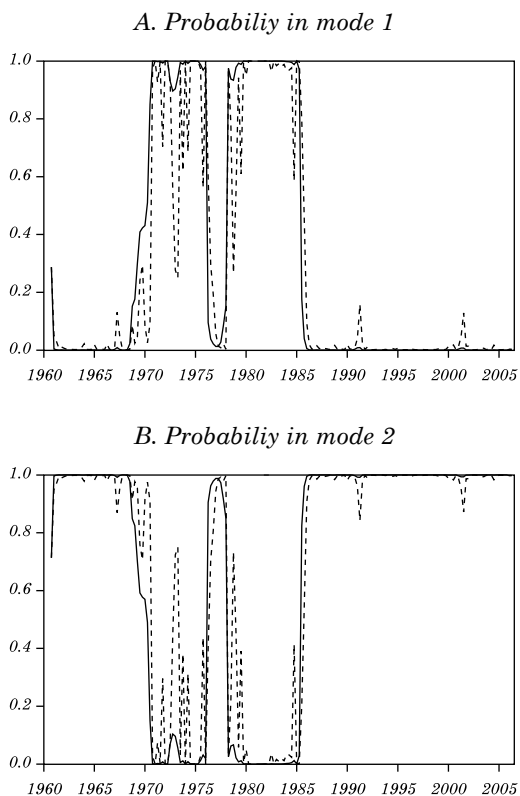
Table 1. Estimates of the Constant-Coefficient Model and a Restricted Two-Mode Lindé Model

<i>Parameter</i>	<i>Mean</i>	<i>Mode 1</i>	<i>Mode 2</i>
ω_f	0.0938	0.3272	0.0000
γ	0.0474	0.0580	0.0432
β_f	0.1375	0.4801	0.0000
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

Source: Authors' calculations.

Mode 2 is thus 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 5. 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 1961–68 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 of the early 1980s. During periods of relative tranquility, such as the Greenspan era, the backward-looking model fits the data the best.

Figure 5. Estimated Probabilities of Being the Different Modes^a



Source: Authors' calculations.

a. In the figure, solid lines graph the smoothed (full-sample) inference, while dashed lines represent the filtered (one-sided) inference.

3.2 Optimal Policy: No Learning and Adaptive Optimal Policy

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

$$i_t = F_i(\mathbf{p}_{t|t})\tilde{\mathbf{X}}_t,$$

where now $\tilde{\mathbf{X}}_t \equiv (\pi_{t-1}, y_{t-1}, y_{t-2}, i_{t-1}, \Xi_{\pi,t-1}, \Xi_{y,t-1})'$. In Svensson and Williams (2007b), we focus on the observable and no-learning cases, and we assume that the shocks $\varepsilon_{\pi t}$ and $\varepsilon_{y t}$ are observable. We thus set $C_2 \equiv 0$ and treat the shocks as additional predetermined variables. 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:

$$L_t = \pi_t^2 + \lambda y_t^2 + \nu (i_t - i_{t-1})^2, \tag{40}$$


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 6. We consider 10,000 simulations of fifty periods, and we plot the median responses for the optimal policy under NL and AOP and the corresponding optimal responses for the constant-coefficient model.¹⁷

Compared with 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 with the constant-coefficient case, the responses of inflation and the output gap are larger and more sustained when there is model uncertainty.

17. 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 6. Unconditional Impulse Responses to Shocks under the Optimal Policy for the Two-Mode Version of the Lindé Model^a



Source: Authors' calculations.

a. In the figure, solid lines represent the median responses under AOP, dashed lines represent the median responses under NL, and dot-dashed lines represent the constant-coefficient responses.

Learning can be beneficial, however, as the optimal policy under AOP dampens the responses to shocks, particularly for shocks to inflation. Since 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 of learning, as by definition they simply provide the responses to a single shock. To gain a better understanding of the role of learning, we simulated our model under the NL and AOP policies to compare the realized economic performance. Table 2 summarizes various statistics resulting from a thousand simulations of a thousand periods each. Thus, for example, the entry for the average π_t is the average across the thousand simulations of the sample average (over the thousand periods) of inflation, while the standard deviation of π_t is the average across simulations of the standard deviation (in each time series) of inflation. In the table, the average period loss (L_t) under AOP is less than half that under NL. Figure 7 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 thousand simulations. The figure shows that the distribution of sample losses is much more favorable under AOP than under NL.

In figure 8 we show one representative simulation to illustrate the differences. The figure reveals that the stabilization of inflation and the output gap are more effective under AOP than NL for very similar instrument rate settings.

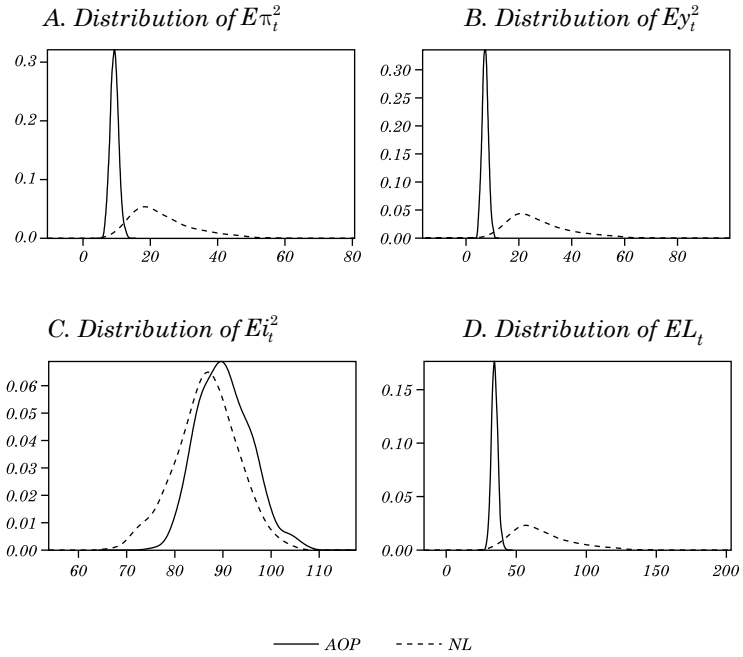
Table 2. Average of Different Statistics under No Learning and the Adaptive Optimal Policy^a

<i>Policy</i>	π_t		y_t		i_t		L_t
	<i>Average</i>	<i>Std. dev.</i>	<i>Average</i>	<i>Std. dev.</i>	<i>Average</i>	<i>Std. dev.</i>	<i>Average</i>
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

Source: Authors' calculations.

a. The table presents the average value and average standard deviation of each variable from one thousand simulations of one thousand periods each of our estimated model under the no learning (NL) and under the adaptive optimal policy (AOP).


Figure 7. Distribution across Samples of Various Statistics under the Optimal Policy for the Two-Mode Version of the Lindé Model



Source: Authors' calculations.

Figure 8. Simulated Time Series under the Optimal Policy for the Two-Mode Version of the Lindé Model^a

A. Inflation



B. Output gap






Figure 8. (continued)

C. Interest rate



D. Probability in mode 1



Source: Authors' calculations.

a. In panels A, B, and C, solid lines denote AOP, while dashed lines graph NL. In panel D, the solid line represents the probability of mode 1, the dotted line represents the true mode, and the dashed gray line represents the unconditional probability of mode 1.

4. 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 has been done to date 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, yet remains relatively tractable in implementation. This is definitely true when decisionmakers 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), which takes the experimentation motive 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.

An issue of much practical importance is the size of the experimentation component of policy and the losses entailed in abstracting from it. While our results in this paper are far from comprehensive, they suggest that the experimentation motive may not be a concern in practical settings. 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. More simulations and cases need to be examined for this to truly be a robust implication.

REFERENCES

- Beck, G.W. and V. Wieland. 2002. "Learning and Control in a Changing Economic Environment." *Journal of Economic Dynamics and Control* 26(9–10): 1359–77.
- Blake, A.P. and F. Zampolli. 2006. "Time Consistent Policy in Markov Switching Models." Working paper 298. London: Bank of England.
- Cogley, T., R. Colacito and T.J. Sargent. 2007. "Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas." *Journal of Money, Credit, and Banking* 39(s1): 67–99.
- Costa, O.L.V., M.D. Fragoso and R.P. Marques. 2005. *Discrete-Time Markov Jump Linear Systems*. London: Springer.
- Do Val, J.B.R. and T. Basar. 1999. "Receding Horizon Control of Jump Linear Systems and a Macroeconomic Policy Problem." *Journal of Economic Dynamics and Control* 23(8): 1099–131.
- Ellison, M. 2006. "The Learning Cost of Interest Rate Reversals." *Journal of Monetary Economics* 53(8): 1895–907.
- Ellison, M. and N. Valla. 2001. "Learning, Uncertainty, and Central Bank Activism in an Economy with Strategic Interactions." *Journal of Monetary Economics* 48(1): 153–71.
- Evans, G. and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Kiefer, N.M. 1989. "A Value Function Arising in the Economics of Information." *Journal of Economic Dynamics and Control* 13(2): 201–23.
- Lindé, J. 2005. "Estimating New-Keynesian Phillips Curves: A Full Information Maximum Likelihood Approach." *Journal of Monetary Economics* 52(6): 1135–49.
- Marcet, A. and R. Marimon. 1998. "Recursive Contracts." Economics working paper 337. Universitat Pompeu Fabra, Department of Economics and Business.
- Rudebusch, G.D. and L.E.O. Svensson. 1999. "Policy Rules for Inflation Targeting." In *Monetary Policy Rules*, edited by J.B. Taylor. University of Chicago Press.
- Svensson, L.E.O. 2007. "Optimization under Commitment and Discretion, the Recursive Saddlepoint Method, and Targeting Rules and Instrument Rules: Lecture Notes." Available at www.princeton.edu/svensson.

- Svensson, L.E.O. and N. Williams. 2007a. "Bayesian and Adaptive Optimal Policy Under Model Uncertainty." Working paper 13414. Cambridge, Mass.: National Bureau of Economic Research.
- . 2007b. "Monetary Policy with Model Uncertainty: Distribution Forecast Targeting." Discussion paper 6331. London: Centre for Economic Policy Research.
- Tesfaselassie, M.F., E. Schaling and S.C.W. Eijffinger. 2006. "Learning about the Term Structure and Optimal Rules for Inflation Targeting." Discussion paper 5896. London: Centre for Economic Policy Research.
- Wieland, V. 2000. "Learning by Doing and the Value of Optimal Experimentation." *Journal of Economic Dynamics and Control* 24(4): 501–34.
- . 2006. "Monetary Policy and Uncertainty about the Natural Unemployment Rate: Brainard-Style Conservatism versus Experimental Activism." *Advances in Macroeconomics* 6(1): 1288–97.
- Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- Zampolli, F. 2006. "Optimal Monetary Policy in a Regime-Switching Economy: The Response to Abrupt Shifts in Exchange-Rate Dynamics." Working paper 297. London: Bank of England.