

Optimal Monetary Policy under Uncertainty: A Markov Jump-Linear-Quadratic Approach*

Lars E.O. Svensson
Sveriges Riksbank and Princeton University
www.princeton.edu/~svensson

Noah Williams
Princeton University
www.princeton.edu/~noahw

March 2008

Abstract

We study the design of optimal monetary policy under uncertainty using a Markov jump-linear-quadratic (MJLQ) approach. We approximate the uncertainty that policymakers face 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 analyze the effects of uncertainty and potential gains from experimentation for two sources of uncertainty in the New Keynesian Phillips curve. Our examples highlight that learning may have sizeable effects on losses, and while it is generally beneficial it need not always be so. The experimentation component typically has little effect, and in some cases it can lead to attenuation of policy.

JEL Classification: E42, E52, E58

Keywords: Optimal monetary policy, learning, recursive saddlepoint method

*The authors thank James Bullard, Timothy Cogley, Andrew Levin, and William Poole 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 National Science Foundation is gratefully acknowledged.

1 Introduction

I have long been interested in the analysis of monetary policy under uncertainty. The problems arise from what we do *not* know; we must deal with the uncertainty from the base of what we *do* know. [...]

The Fed faces many uncertainties, and must adjust its one policy instrument to navigate as best it can this sea of uncertainty. Our fundamental principle is that we must use that one policy instrument to achieve long-run price stability. [...]

My bottom line is that market participants should concentrate on the fundamentals. If the bond traders can get it right, they'll do most of the stabilization work for us, and we at the Fed can sit back and enjoy life.

William Poole [14]

Earlier in his tenure as President of the Federal Reserve Bank of St. Louis, William Poole laid out some of the issues that policymakers face when deciding on policy, as reflected in the quotations above. In this paper we take up some of these issues, applying a framework to help policymakers navigate the “sea of uncertainty.” We focus particularly on the issue of the knowledge and beliefs of the policymakers and the private sector — showing how both groups of agents learn from their observations, and how this may or may not lead to enhanced economic stability. We also address the extent to which policymakers should “sit back,” or instead whether they should actively intervene in markets in order to gain knowledge to help mitigate future uncertainty.

In previous work, Svensson and Williams [17] and [18], 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 in a simple benchmark policy model.

In a first paper, Svensson and Williams [17], 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 learn-

ing (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 [18], 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 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 examples which highlight the effects of learning and experimentation for two sources of uncertainty in the benchmark New Keynesian Phillips curve. 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 ad-

¹ In addition to the classic literature (on such problems as a monopolist learning its demand curve), Wieland [22]-[23] 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 [20] 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.

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.

ditional gains from experimentation.⁴ We find that the experimentation component is typically small. Recognizing the informational component of policy actions often leads policy to be slightly more aggressive, but somewhat surprisingly in one example here it leads to a less aggressive optimal policy.

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. Section 4 presents some conclusions and suggestions for further work.

2 MJLQ Analysis of Optimal Policy

This section summarizes our earlier work, Svensson and Williams [17] and [18].

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

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 [25] has used such an MJLQ model to examine monetary policy under shifts between regimes with and without an asset-market bubble. Blake and Zampolli [2] 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 [17] 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 [15]) and an estimated forward-looking model (a three-mode variant of Lindé [11]). Svensson and Williams [17] 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 [18] focus on learning and experimentation in the MJLQ framework.

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

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, 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 $p_t = (p_{1t}, \dots, p_{n_jt})'$ 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

⁶ See also Svensson and Williams [17], where we show how many different types of uncertainty can be mapped into our MJLQ framework.

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.

and the determination of $p_{t|t}$ will be specified below.

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

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

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

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

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 [17], 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

$$p_{t+1|t+1} = P'p_{t|t}. \quad (2.7)$$

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 [17].

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 [17] 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,

$$\mathbf{E}_t H_{j_{t+1}} x_{t+1} = z_t, \quad (2.8)$$

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

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

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

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 Marcat and Marimon [12], Svensson and Williams [17], and Svensson [16] 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 [17], 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

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

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

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

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 [17] and [18] 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

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

⁹ To be precise, the observation of X_t , which depends on $C_{1j_t \varepsilon_t}$, allows some inference of $\varepsilon_t, \varepsilon_{t|t}$. 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 $E_t C_{2j_t \varepsilon_t} = 0$.

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 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

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

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 .

¹⁰ The policymaker and private sector can also estimate the shocks $\varepsilon_{t|t}$ as $\varepsilon_{t|t} = \sum_j p_{jt|t} \varepsilon_{jt|t}$, where $\varepsilon_{jt|t} \equiv X_t - A_{11j} X_{t-1} - A_{12j} x_{t-1} - B_{1j} i_{t-1}$ ($j \in N_j$). 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 $\varepsilon_{jt|t}$.

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). 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

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

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).

Eight, 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

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

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 $x(s_t, j_t, 0_t)$.

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 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 [18].

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, $\check{V}(s_t, j_t)$, satisfies

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

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 [18], 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 \tilde{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.

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

The transition equation for the BOP case is:

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

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 right side of (2.12) can be written as

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

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

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

See Svensson and Williams [18] for a proof.

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

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

where the transition equation is given by (2.30).

The solution to the optimization problem can be written

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

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

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.

2.7 Observable Modes

In this paper we largely focus on the cases where the policymakers do not observe the current mode, which is certainly the more relevant case when analyzing model uncertainty. However some situations may arguably be better modeled by observable shifts in modes, as in most of the econometric literature on regime switching models. Moreover one way to gauge the effects of uncertainty in the a model is to move from a constant-coefficient specification to one in which the parameters are observable but may vary. (That is, the current values of parameters are known, but future values are uncertain.) We use the observable mode case for this reason, to analyze implications of uncertainty on policy. In Svensson and Williams [17] we develop simple algorithms for observable changes in modes, which play off the fact that conditional on the mode the evolution of the economy is linear and preferences are quadratic. Thus the optimal policy consists of a mode-dependent collection of linear policy rules and can be written

$$i_t = F_{i_{j_t}} \tilde{X}_t. \quad (2.35)$$

for $j_t \in N_j$.

3 Learning and experimentation in a simple New Keynesian model

3.1 The model

For our policy exercises, we consider a benchmark hybrid New Keynesian Phillips curve (see Woodford [24] for an exposition):

$$\pi_t = (1 - \omega_{j_t}) \pi_{t-1} + \omega_{j_t} \mathbf{E}_t \pi_{t+1} + \gamma_{j_t} y_t + \varepsilon_t \quad (3.1)$$

Here π_t is the inflation rate, y_t is the output gap, ω_{j_t} is a parameter reflecting the degree of forward-looking behavior in price setting, and γ_{j_t} is a composite parameter reflecting the elasticity of demand and frequency of price adjustment. For simplicity, we assume that policymakers can directly control the output gap y_t . In another paper, Svensson and Williams [19] we consider optimal policy in the standard two equation New Keynesian model which also includes a log-linearized consumption Euler equation. Many of the same issues that we focus on here arise there as well, but the simpler setting in the present paper allows us to focus more directly on the effects of uncertainty on policy.

We focus on two key sources of uncertainty in the New Keynesian Phillips curve. Our first example considers the degree of forward-looking behavior in inflation. In the model, this translates to uncertainty about ω_j . If this parameter is large, inflation is largely determined by current shocks and expectations of the future, while if ω_j is small then there is a substantial exogenous inertia in the inflation process. Our second example analyzes uncertainty about the slope of the Phillips curve, as reflected in the parameter γ_j . 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. In each example we look first at the effect of uncertainty, going from a constant-coefficient model to a model with random coefficients. Then we analyze the effects of learning and experimentation on policy and losses.

In both examples we use the following loss function,


$$L_t = \pi_t^2 + \lambda y_t^2, \tag{3.2}$$

We set the loss function parameters as: $\delta = 0.98$, $\lambda = 0.1$, and set the shock standard deviation to $c = 0.5$. Even though different structural parameters vary in the two examples, both examples use two modes and set the transition matrix to:

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

In both examples we examine the 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 have one forward-looking variables ($x_t \equiv \pi_t$) and consequently one Lagrange multiplier ($\Xi_{t-1} \equiv \Xi_{\pi,t-1}$). We have one predetermined variable ($X_t \equiv \pi_{t-1}$) 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 three dimensional ($s_t = (\pi_{t-1}, \Xi_{\pi,t-1}, p_{1|t})'$). For computational reasons we are forced to restrict attention to relatively sparse grids with few points. The following

Figure 3.1: Policies and losses from observable (Obs) and constant modes



plots show two-dimensional slices of the value and policy functions, focusing on the dependence on π_{t-1} and $p_{1t|t}$ (which we for simplicity denote by p_{1t} in the figures). In particular, all of the plots are for $\Xi_{\pi,t-1} = 0$.

3.2 Example 1: How forward-looking is inflation?

This example analyzes one of the main sources of uncertainty in the New Keynesian framework, the degree to which inflation is a forward-looking variable responding to expectations of future developments. Specifications which suggest that inflation has substantial exogenous persistence have tended to fit better empirically, while perhaps being less rigorous in their micro-foundations. In this example we see how uncertainty about the degree of forward-looking behavior, as indexed by ω_j , affects policy. Thus we assume that there are two modes, one more forward-looking which has $\omega_1 = 0.8$ and the other more backward-looking with $\omega_2 = 0.2$. Note that with the transition matrix P as specified above, this means $E(\omega_j) = 0.5$. For this example, we fix the slope parameter at $\gamma = 0.1$.

In figure 3.1 we illustrate the effects of uncertainty on policy and losses. In the left panel, we plot the two mode-dependent optimal policy functions for the MJLQ model with observable modes, labeled “Obs 1” for mode 1 and “Obs 2” for mode 2. Here we see that the optimal policy is more aggressive in the more backward-looking mode 2, as in response to a higher inflation the optimal policy involves larger negative output gaps. The unconditional average policy is labeled “E(Obs)” and shown with a solid line. For comparison, we also plot with a dotted line the constant-coefficient

case where we set $\omega_1 = \omega_2 = E(\omega_j) = 0.5$. Here we see that the optimal policy under uncertainty is more aggressive in responding to inflation movements than in the absence of uncertainty.

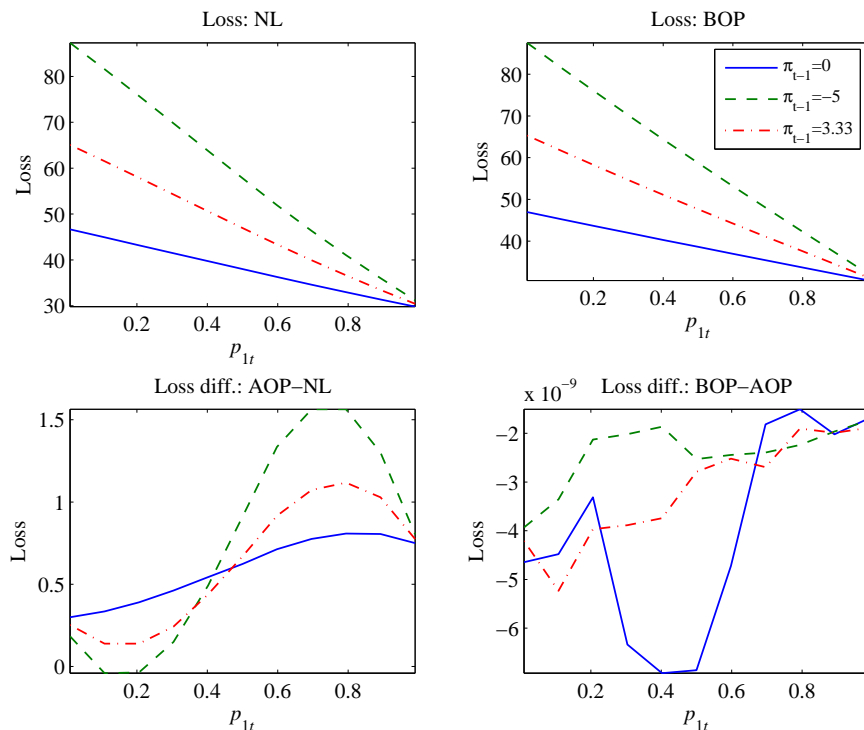
A common starting point for thinking about the effects of uncertainty on policy is Brainard's [3] classic analysis which suggested that uncertainty should make policy more cautious. However Brainard worked in a static framework, and the source of uncertainty he analyzed was a slope coefficient on how policy affects the economy. Our second example below is closer to Brainard's, and comes to similar conclusions. But in this example our results suggest, at least for this parameterization, that uncertainty about the dynamics of inflation leads to more aggressive policy. This is similar to what Söderström [21] found in a backward-looking model.

The right panel of figure 3.1 plots the losses associated with the optimal policies in the different cases. When inflation is more forward-looking, it is easier to control and so overall losses are lower even with less aggressive policies. However uncertainty about the dynamics of inflation can have significant effects on losses for moderate to high inflation levels. This is evident by comparing the constant-coefficient and average observable curves, where we see that the loss nearly doubles at the edges of the plot.

Now we keep the same specification, but make the more realistic assumption that the current mode is not observed. Thus we analyze the effects of learning and experimentation on policy and losses. The top two panels of figure 3.2 show losses under NL and BOP as functions of p_{1t} . The bottom two panels of the figure show the differences between losses under NL, AOP, and BOP. Figure 3.3 shows the corresponding policy functions and their differences. The top two panels plot the policy functions under AOP and BOP as a function of inflation. The AOP policy is linear in π_t , and clearly the BOP policy is nearly so. The bottom left panel plots the BOP policy as a function of p_{1t} , showing that policy is less aggressive (that is, smaller magnitude of response) the greater is the probability of being in the more forward-looking mode 1. The bottom right panel shows that the policy differences between AOP and BOP, the experimentation component of policy, are incredibly small.

In Svensson and Williams [18] 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 3.2 that the value function is actually slightly convex in p_{1t} , so learning is not beneficial here. Consequently, we see in figure 3.2 that AOP gives higher losses than NL. In contrast, for a backward-looking

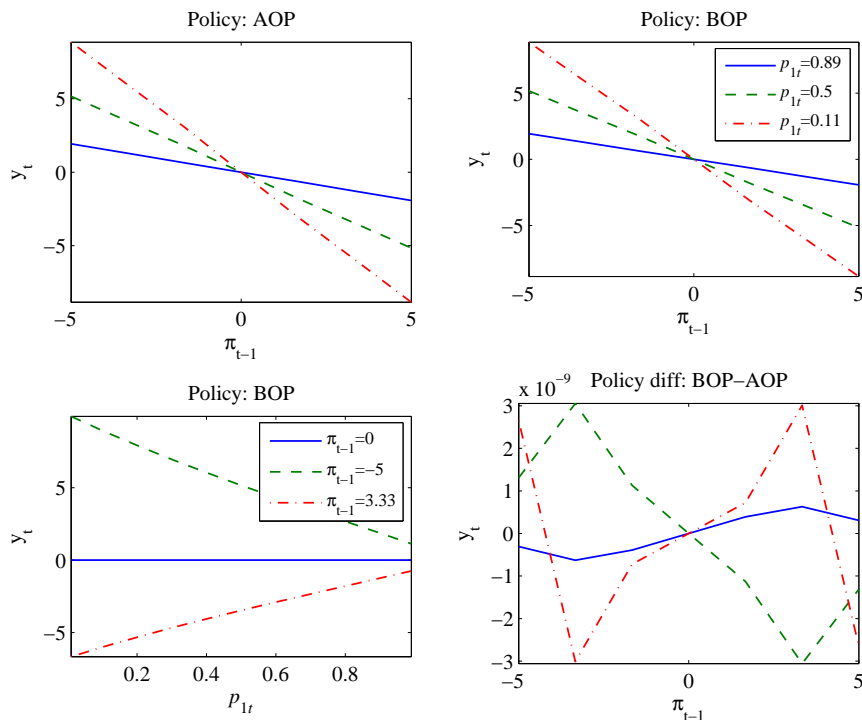
Figure 3.2: Losses and differences in losses from no learning (NL), adaptive optimal policy (AOP) and Bayesian optimal policy (BOP)



example in Svensson and Williams [18], the value function is concave and learning is beneficial. Experimentation is beneficial here, as BOP does give lower losses than AOP, but the difference is miniscule. So for this example, learning has sizeable effects on losses and is detrimental, while experimentation is beneficial but has negligible effects.

Why would learning not be beneficial 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. However 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 between backward- and forward-looking models then comes from the way that private sector beliefs also respond to learning. Having more reactive private sector beliefs may add volatility and make it more difficult for the policymaker to stabilize the economy.

Figure 3.3: Optimal policies and their differences under adaptive optimal policy (AOP) and Bayesian optimal policy (BOP)

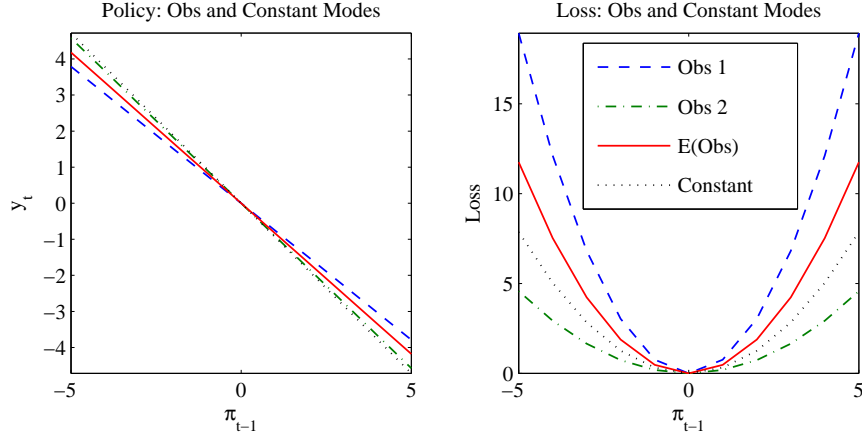


3.3 Example 2: What is the slope of the Phillips curve?

This example analyzes the other main source of uncertainty in the New Keynesian Phillips curve, the extent to which inflation responds to fluctuations in the output gap. Once again we assume that there are two modes, now one in which the Phillips curve is flatter which has $\gamma_1 = 0.05$ and the other has a steeper curve with $\gamma_2 = 0.25$. Note that with the transition matrix P as specified above, this means $E(\gamma_j) = 0.15$. For this example, we fix the forward-looking expectations parameter at $\omega = 0.5$. Since policymakers once again directly control the output gap, this example is a forward-looking counterpart to the classic Brainard [3] analysis of uncertainty about the effectiveness of the control.

In figure 3.4 we illustrate the effects of uncertainty on policy and losses. As in the previous example, the left panel plots the two mode-dependent optimal policy functions for the MJLQ model with observable modes. Here we see that the MJLQ optimal policies in both modes are less aggressive than the constant-coefficient case. Thus our results here are in accord with Brainard’s — uncertainty about the slope of the Phillips curve leads to more cautious policy.

Figure 3.4: Policies and losses from observable (Obs) and constant modes




The right panel of figure 3.4 plots the losses associated with the optimal policies in the different cases. When the Phillips curve is steeper, inflation responds more to the output gap making it easier to control. Thus overall losses are lower in mode 2 even with less aggressive policies. However once again uncertainty about this key parameter can have significant effects on losses for high inflation levels. This is evident by comparing the constant-coefficient and average observable curves, where we see that the loss nearly doubles at the edges of the plot.

Now we again keep the same specification, but make the more realistic assumption that the current mode is not observed. The top two panels of figure 3.5 show losses under NL and BOP as functions of p_{1t} . The bottom two panels of the figure show the differences between losses under NL, AOP, and BOP. We see in figure 3.2 that the value function is once again slightly convex in p_{1t} , so learning is not beneficial here. Consequently, we see in the bottom right panel of figure 3.2 that AOP gives higher losses than NL. Thus once again, the additional volatility outweighs the improved inference, and makes learning detrimental in this example. Experimentation is once again beneficial, as BOP gives lower losses than AOP. And while the effects of experimentation are an order of magnitude smaller than the effects of learning, the gains from recognizing the endogeneity of information are non-negligible here. Thus for uncertainty about the slope of the Phillips curve, policymakers may have an incentive to experiment — that is, to take actions to mitigate future uncertainty.

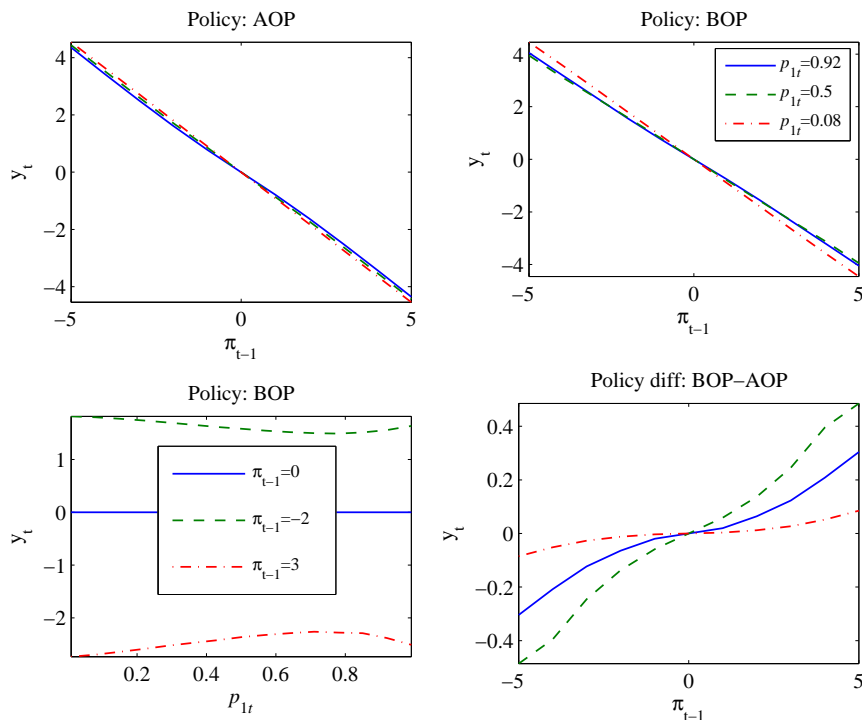
Figure 3.6 shows the corresponding policy functions and their differences. The top two panels plot the policy functions under AOP and BOP as a function of inflation. The AOP policy is linear

Figure 3.5: Losses and differences in losses from no learning (NL), adaptive optimal policy (AOP) and Bayesian optimal policy (BOP)



in π_t , and clearly the BOP policy is nearly so, although some differences are evident at the edge of the plot. The bottom left panel plots the BOP policy as a function of p_{1t} , showing that the policy function is relatively flat in this dimension. The bottom right panel plot the difference between the AOP and BOP policy functions, which shows that here the experimentation motive leads toward less aggressive policy. This is counter to an example in Svensson and Williams [18], who show that in a backward-looking model experimentation may lead to more aggressive policy. There policy makes outcomes more dispersed in order to sharpen inference over the modes. However here because learning is detrimental, the experimentation component of policy seeks to slow the effects of learning by making outcomes less dispersed. This serves to illustrate that the experimentation component of policy need not be associated with wild or aggressive policy action, but rather it optimally takes into account how information influences the targets of policy.

Figure 3.6: Optimal policies and their differences under adaptive optimal policy (AOP) and Bayesian optimal policy (BOP)



4 Conclusion

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] Blake, Andrew P., and Fabrizio Zampolli (2006), “Time Consistent Policy in Markov Switching Models,” Working Paper No. 298, Bank of England.
- [3] Brainard, W. (1967), “Uncertainty and the Effectiveness of Policy,” *American Economic Review* 57, 411–425.
- [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., Marcelo 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] Marcat, 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] Poole, William (1998), “A Policymaker Confronts Uncertainty,” *Federal Reserve Bank of St. Louis Review*, September 1998, 3–8.
- [15] 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.
- [16] 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.
- [17] 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/.
- [18] 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/.
- [19] Svensson, Lars E.O., and Noah Williams (2008), “Optimal Monetary Policy in DSGE Models: A Markov Jump-Linear-Quadratic Approach,” working paper, www.princeton.edu/svensson/ and www.princeton.edu/~noahw/.

- [20] 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.
- [21] Söderström, Ulf (2002), “Monetary Policy with Uncertain Parameters,” *Scandinavian Journal of Economics*, 104, 125-145.
- [22] Wieland, Volker (2000), “Learning by Doing and the Value of Optimal Experimentation,” *Journal of Economic Dynamics and Control* 24, 501–534.
- [23] Wieland, Volker (2006), “Monetary Policy and Uncertainty about the Natural Unemployment Rate: Brainard-Style Conservatism versus Experimental Activism,” *Advances in Macroeconomics* 6(1).
- [24] Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.
- [25] 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.