Monetary Policy under Financial Uncertainty☆

Noah Williams

University of Wisconsin - Madison

Abstract

Monetary policy may play a substantial role in mitigating the effects of financial crises. In this paper, I suppose that the economy occasionally but infrequently experiences crises, where financial variables affect the broader economy. I analyze optimal monetary policy under such financial uncertainty, where policymakers recognize the possibility of crises. Optimal monetary policy is affected during the crisis and in normal times, as policymakers guard against the possibility of crises. In the estimated model this effect is quite small. Optimal policy does change substantially during a crisis, but uncertainty about crises has relatively little effect.

Keywords: Optimal monetary policy, financial crises, model uncertainty

JEL Classification: E42, E52, E58

1. Introduction

The recent financial crisis and subsequent recession have illustrated how developments in credit and financial markets may be transmitted to the economy as a whole. However prior to the crisis, the baseline models for monetary policy analysis had no direct way to model such developments. The potential importance of financial factors was recognized in the literature, but financial factors were not present in the most widely-used models for policy analysis. One interpretation of this state of affairs is that in "normal times" financial market conditions are not of primary importance for monetary policy. In such times, policy focuses on the consequences of interest rate setting for inflation and output, reacting primarily to shocks which directly affect these variables. However the economy may occasionally enter "crisis" periods when financial frictions are of prime importance and shocks initially affecting financial markets may in turn impact the broader economy. The transitions between normal

☆I thank Lars E.O. Svensson, without implicating him in the faults of this paper. Our joint work forms the basis of all the analysis here. I thank Marvin Goodfriend and Chris Sims for helpful comments.

Preprint submitted to Elsevier *April 13, 2012*

and crisis period may be difficult to predict, and a crisis may be well underway before its effects become apparent in the broader economy. In this paper I develop methods to provide guidance in assessing and responding to such financial uncertainty.

In this paper, I focus on monetary policy design when occasional crisis episodes impact on the transmission mechanism. Importantly, we do not consider financial stability policy, which may have distinct objectives (financial stability, appropriately defined) and instruments (bank supervision and regulation, liquidity provision to banks, and so on). In our setting, monetary policy always has as its objective the stabilization of inflation around a target and economic activity around a target of a sustainable level, and sets a nominal interest rate as its instrument. Crises impact the ability of monetary policymakers to attain these objectives, as they introduce additional shocks and factors which affect inflation and output. Importantly, we take crises here as exogenous, reflecting financial market developments beyond the control of monetary policy. Thus we focus on how monetary policy may mitigate the effects of such crises, and how uncertainty about financial crises affects the appropriate monetary policy response.

This paper encapsulates a stylized reading of the developments in monetary policy analysis over the past decade. By the mid-2000s there had been influential work showing that larger New Keynesian models were able to successfully confront the aggregate data. In particular, the work of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2003) showed that such theoretically-based models were able to fit aspects of the data comparable to VARs. Such models incorporated a host of real and nominal frictions, but did not discuss financial factors. In addition, there was a growing literature on monetary policy analysis under uncertainty, some of which used these larger scale models.[1] This literature considered the implications for policy of model uncertainty, including uncertainty about the specifications and parameterizations of the models, and the types of nominal rigidities. But again financial factors were notably (in hindsight) absent. Of course, the seminal contributions of Bernanke and Gertler (1989), Kiyotaki and Moore (1997), and Bernanke, Gertler,

[1] A very brief and highly selective list of references includes work by Onatski and Stock (2002), Giannoni (2002), Levin, Wieland, and Williams (2003), and Levin, Onatski, Williams, and Williams (2006).

and Glichrist (1999) were recognized. There was also ongoing work on financial frictions in monetary policy, including work by Christiano, Motto, and Rostagno (2003) and Gertler, Gilchrist, and Natalucci (2007) among others. But the "consensus" policy models had not yet incorporated these frictions. The turmoil of the past several years has naturally spurred interest in models of financial frictions and the interaction of real and financial markets more broadly.

In hindsight, it is clear that the much of the previous literature on monetary policy analysis missed a big source of uncertainty: uncertainty about financial sector impacts on the broader economy. Under one reading, this was simply an omission, and monetary policymakers should have been more focused on financial factors throughout. In this paper we suggest another interpretation, namely that there may be significant variation over time in the importance of financial shocks for monetary policy. In normal times, defaults and bank failures are rare, sufficient liquidity is provided for businesses, and monetary policy focuses on responding to shocks to inflation and output. However in crisis periods, defaults and bank failures increase, liquidity may be scarce, and shocks to the financial sector may impact the transmission of monetary policy. I assume that the economy switches stochastically between such "normal times" and "crisis" regimes, and consider the design of monetary policy in an environment where policymakers and private sector agents recognize the possibility of such switches.

As a model of "normal times" I use a small empirical New Keynesian model. In particular, I use a version of the model of Lindé (2005), which adds some additional exogenous persistence in the form of lagged dynamics to the standard New Keynesian model. For the model of crises, I use a version of the model of Curdia and Woodford (2009b), which is a tractable extension of the standard New Keynesian model to incorporate financial frictions. As in the standard model, the key equilibrium conditions of the model include a log-linearized consumption Euler equation (governing aggregate demand) and a New Keynesian Phillips curve (reflecting price setting with nominal rigidities). However the allocative distortions associated with imperfect financial intermediation give rise to a spread between borrowing and lending interest rates, and a gap in the marginal utility between borrowers and lenders. These factors only matter for inflation and output determination in a crisis, and an exogenous Markov chain governs the switches of the economy between normal and crisis periods. Importantly, I focus on a simple specification of the model where the key interest rate spread is exogenous.

I first suppose that crises are observable, so the main source of uncertainty is over the future state of the economy. I then consider the case where agents must infer the current state of the economy from their observations, so uncertainty and learning about the current state become additional considerations. Thus even in normal times, the optimal policy differs from the prescriptions of a model without such crises. The optimal policy under uncertainty reflects the possibility that the economy may transit into a crisis in the future, as well as the uncertainty about whether the economy may already have switched into such a state. Thus the results imply variation over time in the policy response to shocks to real and financial factors, with learning about the state of the economy potentially playing a role in moderating fluctuations.

The policy analysis in this uses the approach of Svensson and Williams (2007b) and (2007a). There 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.[2] 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 I put those suggestions into practice, by analyzing uncertainty about financial factors and the transmission of financial shocks to the rest of the economy.

In a first paper, Svensson and Williams (2007b), we studied optimal policy design in MJLQ models when policymakers can or cannot observe the current mode, but we abstracted from any learning and inference about the current mode. Although in many cases the optimal policy under no learning (NL) is not a normatively desirable policy, it serves as a

[2]Related approaches are developed by Blake and Zampolli (2006), Tesfaselassie, Schaling, and Eijffinger (2006), Ellison and Valla (2001), Cogley, Colacito, and Sargent (2007), and Ellison (2006).

useful benchmark for our later policy analysis. In a second paper, Svensson and Williams (2007a), 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. 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).[3] 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 (2001) for an overview). Further, the AOP specification rules out the experimentation which some may view as objectionable in a policy context.[4] In this paper, I apply our methodology to study optimal monetary-policy design under what I call "financial uncertainty."

Overall, I find that in the estimated model the optimal monetary policy does change substantially during a crisis, but uncertainty about crises has relatively little effect. In crises, it is optimal for the central bank to cut interest rates substantially in response to increases in the interest rate spread. However the size of this response is nearly the same in our MJLQ model as in the corresponding constant coefficient model. In addition, the possibility that the economy may enter a crisis means that even in normal times policy should respond to interest rate spreads. But again, this effect is fairly negligible. These

[3] 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.

[4] 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.

results seem to rely on two key factors: the exogeneity of the interest rate spreads and the rarity of crises. In regard to the first point, policy cannot affect spreads in our model, so responding to interest rate spreads in normal times has no effect on the severity of crises. If policy could affect spreads, then there may be more of a motive for policy to react before a crisis would appear, as stabilizing interest spreads may make crises less severe. On the second point, note that by responding to spreads in normal times policymakers are effectively trading off current performance for future performance. The greater the chance of transiting into a crisis, the larger the weight that the uncertain future would receive in this tradeoff. As crises are sufficiently rare, there is little reason to sacrifice much current performance. Policymakers are typically able to react sufficiently strongly once crises do arrive, so there is little reason to alter policy in advance of the crisis.

Our conclusions are certainly model-specific, and as we've noted, they rely on the exogeneity of interest rate spreads. Certainly during the crisis most central banks rapidly expanded their balance sheets, making asset purchases as a means of providing liquidity to financial markets and attempting to reduce interest rate spreads. In this paper I focus on interest rate policy solely, treating liquidity policy as a separate issue. Curdia and Woodford (2009a) show that in their model, as used in this paper, liquidity policy can indeed be viewed as a separate instrument which need not affect interest rate policy. But in general there may be broader interactions, with liquidity policy imposing costs, such as political pressure associated with the central bank holding a broader array of assets, which could affect future interest rate policy. Such issues are clearly relevant for the current policy environment, but are outside the scope of this paper.

The paper is organized as follows: Section 2 presents the MJLQ framework and summarizes our earlier work. Section 3 then develops and estimates our benchmark model of financial uncertainty, while Section 4 analyzes optimal policy in the context of this model under different informational assumptions. Section 5 presents some conclusions and suggestions for further work.

2. MJLQ Analysis of Optimal Policy

This section summarizes our earlier work, Svensson and Williams (2007b) and (2007a). Here we outline the approach that we use to structure and analyze uncertainty in this paper.

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 $j_t \in N_j \equiv \{1, 2, ..., n_j\}$ denote the mode in period t. The model of the economy can then be written

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

$$\mathrm{E}_t H_{j_{t+1}} x_{t+1} = A_{21j_t} X_t + A_{22j_t} x_t + B_{2j_t} i_t + C_{2j_t} \varepsilon_t, \tag{2}$$

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

Note that the matrices on the right side of (1) depend on the mode j_{t+1} in period $t+1$, whereas the matrices on the right side of (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 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 $j \in N_j$.

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

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

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., and are the driving forces in the model. They may not be directly observed. It is convenient but not necessary that they are independent of each other and the mode. We let $p_t = (p_{1t}, ..., p_{n_j t})'$ denote the true probability distribution of j_t in period t. We let $p_{t+\tau|t}$ denote the policymaker's and private sector's estimate in the beginning of period t of the probability distribution in period $t+\tau$. The *prediction* equation for the probability distribution is

$$p_{t+1|t} = P' p_{t|t}. \tag{3}$$

We let the operator $\mathrm{E}_t[\cdot]$ in the expression $\mathrm{E}_t H_{j_{t+1}} x_{t+1}$ on the left side of (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 (in general) excluding j_t and ε_t. Thus, we assume that information is symmetric between the policymaker and the (aggregate) private sector. Our methods can be easily adapted to consider information asymmetries as well. Although we focus on the determination of the optimal policy instrument i_t, our results also show how private sector choices as embodied in x_t are affected by uncertainty and learning. The precise informational assumptions and the determination of $p_{t|t}$ will be specified below.

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

$$\mathrm{E}_t \sum_{\tau=0}^{\infty} \delta^\tau L(X_{t+\tau}, x_{t+\tau}, i_{t+\tau}, j_{t+\tau}) \tag{4}$$

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

$$L(X_t, x_t, i_t, j_t) \equiv \begin{bmatrix} X_t \\ x_t \\ i_t \end{bmatrix}' W_{j_t} \begin{bmatrix} X_t \\ x_t \\ i_t \end{bmatrix}, \tag{5}$$

where the matrix W_j ($j \in N_j$) is positive semidefinite.

We assume that the policymaker optimizes under commitment in a timeless perspective, although our methods directly extend to other cases as well. To solve for optimal policies, we use the recursive saddlepoint method of Marcet and Marimon (1998) to extend the methods

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

for MJLQ models developed in the control theory literature to allow for forward looking endogenous variables. We thus supplement the state vector X_t with the vector Ξ_{t-1} of lagged Lagrange multipliers for equation (2). The timeless perspective requires that we then add the term

$$\Xi_{t-1}\frac{1}{\delta}\mathrm{E}_t H_{j_t} x_t \tag{6}$$

to the intertemporal loss function in period t. The current values of the Lagrange multipliers, which we denote γ_t, becomes an additional control vector, and the state vector is supplemented with the additional equation:

$$\Xi_t = \gamma_t.$$

Additionally, the period loss function is supplemented with the Lagrangian terms in the multiplier γ_t and the constraint (2). On this expanded state space, system (1)-(2) can be solved as a MJLQ model, where the objective is minimized with respect to i_t but maximized with respect to (x_t, γ_t).

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 (2007b), and present an illustration. 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 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 highly persistent Markov chains, our MJLQ provide accurate approximations of nonlinear models with Markov switching.

2.3. Types of optimal policies

We will distinguish four cases of optimal policies: (1) Optimal policy when the modes are observable (OBS), (2) Optimal policy when there is no learning (NL), (3) Adaptive optimal policy (AOP), and (4) Bayesian optimal policy (BOP). Here we briefly discuss the different cases, deferring to Svensson and Williams (2007b) and (2007a) for details.

The most direct case is when the policymaker and the private sector directly observe the modes (OBS). This is typically the case studied in the econometric literature on regime switching, where agents implicitly observe the current regime but the econometrician does not. Similar approaches have also been used in the literature on "policy switching". Under OBS, the optimal policy conditions on the current mode, taking into account that the mode may switch in the future. Svensson and Williams (2007b) show that optimal policies in this case consist of mode-dependent linear policy rules, which can be computed efficiently even in large models. The conditionally linear-quadratic structure that the MJLQ approach provides great simplicity in this setting.

The other three cases all suppose that the modes are not observable by the policymakers (and the public). The cases differ in their assumptions about how policymakers use observations to make inferences about the mode, and how they use that information to form policy. 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

$$p_{t+1|t+1} = P' p_{t|t}. \tag{7}$$

That is, the policymaker and the private sector do not use observations of the economy to update the probability distribution. The policymaker then determines optimal policy in period t conditional on $p_{t|t}$ and (7). This is a variant of a case examined in Svensson and Williams (2007b). Since the beliefs evolve exogenously, the tractability of the MJLQ structure is again preserved, and computations are quite simple.

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 economy to update the 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. The AOP case is simple to implement recursively, as we have already discussed how to solve for the optimal decisions, and the Markov structure allows for simple updating of probabilities. However, the ex-ante evaluation of expected loss is more complex, as it must account for the nonlinearity of the belief updating.

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 channel into account. Therefore, BOP includes optimal experimentation, where for instance the policymaker may pursue a policy that increases losses in the short run but improves the inference of the probability distribution and therefore lowers losses in the longer run. Although policymakers sometimes express skepticism about policy experimentation, it is a natural byproduct of optimal policy. In practical terms, the fact that the updating equation for beliefs is nonlinear means that more complex numerical methods are necessary in this case. Practically speaking, computational considerations mean that BOP is only feasible in relatively small models.

As we discuss in Svensson and Williams (2007a), Bayesian updating makes beliefs respond to information, and thus increases their volatility. Thus the curvature of the value function will influence whether learning is beneficial or not. In some cases the losses incurred by increased variability of beliefs may offset the expected precision gains. This may be particularly true in forward-looking models where policymakers and the private sector share the same beliefs. Learning by the private sector may induce more volatility, thus making it more difficult for policymakers to stabilize the economy. We show below how these issues manifest themselves in the applications.

What makes models with forward-looking variables different? One difference is that with backward-looking models, the BOP is always weakly better than the AOP, as acknowledging the endogeneity of information in the BOP case need not mean that policy must change. That is, the AOP policy is always feasible in the BOP problem. However, with forward-looking models, neither of these conclusions holds. Under our assumption of symmetric information and beliefs between the private sector and the policymaker, both the private sector and the policymaker learn. If we allow beliefs to differ, then the BOP is always weakly for policymakers to learn, given private sector behavior. This is just as in the backward-looking case. Forward-looking models differ in 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. With symmetric beliefs, acknowledging the endogeneity of information in the BOP case need not be beneficial, as it may induce further volatility in agents' beliefs.[8]

3. Uncertainty about the impact of financial variables

3.1. Overview

In this section we consider our benchmark formulation of financial uncertainty, where policymakers are uncertain about the impact of financial variables on the broader economy, and show how to incorporate such uncertainty in a MJLQ model. This section implements one of the scenarios outlined in the introduction, that in "normal times" financial market conditions are not important for monetary policy. We capture this assumption by taking one mode of our MJLQ model to be a relatively standard New Keynesian model, in particular a version of Lindé's (2005) empirical model of US monetary policy. However the economy may occasionally enter "crisis" periods when financial market frictions and potential credit market disruptions imply that financial variables may impact the broader economy. We take a direct approach to this, based on the work of Curdia and Woodford (2009b). They develop a modification of the standard New Keynesian model which incorporates a credit spread as an additional factor influencing output and inflation. Thus we assume that in the "crisis" mode credit spreads matter for monetary policy, but in normal times they do not. We then calibrate and estimate the model using recent US data, and analyze the optimal policies under different informational assumptions. We are particularly interested in analyzing not

[8]Technically, these results are manifest in fact that in the forward-looking case we solve saddlepoint problems. So by going from AOP to BOP we are expanding the feasible set for both the minimizing and maximizing choices.

only how the optimal monetary policy differs in crises, but also how the knowledge that crises are possible affects the optimal policy in normal times.

3.2. The model

We now lay out the model in more detail. As discussed above, one mode represents "normal times," via a typical small but empirically plausible model. We consider a variation on the benchmark "three equation" New Keynesian model, consisting of a New Keynesian Phillips curve, a consumption Euler equation, and a monetary policy rule (see Woodford (2003) for an exposition). We focus on a version of the model of Lindé (2005), which we also we estimated in Svensson and Williams (2007b). Compared to the standard New Keynesian model, this model includes richer dynamics for inflation and the output, as both have backward- and forward-looking components. In particular, the model in normal times is given by:

$$\begin{aligned} \pi_t &= \omega_f E_t \pi_{t+1} + (1-\omega_f)\pi_{t-1} + \gamma y_t + c_\pi \varepsilon_{\pi t}, \\ y_t &= \beta_f E_t y_{t+1} + (1-\beta_f)\left[\beta_y y_{t-1} + (1-\beta_y) y_{t-2}\right] - \beta_r \left(i_t - E_t \pi_{t+1}\right) + c_y \varepsilon_{yt}. \end{aligned} \tag{8}$$

Here π_t is the inflation rate, y_t is the output gap, and i_t is the nominal interest rate, and the shocks $\varepsilon_{\pi t}$, ε_{yt} are independent standard normal random variables. For empirical analysis, we supplement the model with flexible Taylor-type policy rule:

$$i_t = (1-\rho_1-\rho_2)\left(\gamma_\pi \pi_t + \gamma_y y_t\right) + \rho_1 i_{t-1} + \rho_2 i_{t-2} + c_i \varepsilon_{it} \tag{9}$$

where the policy shock ε_{it} is also an i.i.d. standard normal random variable.

To this relatively standard depiction of monetary policy in normal times, we now add the possibility of a "crisis" mode, or more precisely, a mode in which credit spreads matter for inflation and output determination. As discussed above, we use a version of the Curdia-Woodford (2009b) model which adds credit market frictions to the standard New Keynesian model. The model results in a spread between borrowing and deposit interest rates (a credit spread), and heterogeneity across borrowers and savers which is reflected in a marginal utility gap between them. We focus on the version of the model where the credit spread is exogenous, although Curdia and Woodford also consider a specification which endogenizes the spread.

The exogeneneity of the spread results in rather stark differences in policy responses across modes, and allows us to focus on the policy response to credit spreads.

In our specification of the crisis mode, we keep the dynamics of the Lindé model, but supplement it with a credit spread ω_t and the marginal utility gap Ω_t between borrowers and savers. Thus the model in crisis times is given by:

$$\begin{aligned}
\pi_t &= \omega_f E_t \pi_{t+1} + (1-\omega_f)\pi_{t-1} + \gamma y_t + \xi \Omega_t + c_\pi \varepsilon_{\pi t}, \qquad (10)\\
y_t &= \beta_f E_t y_{t+1} + (1-\beta_f)\left[\beta_y y_{t-1} + (1-\beta_y) y_{t-2}\right] - \beta_r \left(i_t - E_t \pi_{t+1}\right) + \theta \Omega_t + \phi \omega_t + c_y \varepsilon_{yt}.\\
\Omega_t &= \delta E_t \Omega_{t+1} + \omega_t\\
\omega_{t+1} &= \rho_\omega \omega_t + c_\omega \varepsilon_{\omega t+1}.
\end{aligned}$$

Thus, in addition to the new variables entering the equations for inflation and the output gap, we now have the endogenous dynamics of the marginal utility gap Ω_t as well as the exogenous dynamics of the interest spread ω_t. We assume that the spread follows an AR(1) process, where again the shock to the spread $\varepsilon_{\omega t}$ is an i.i.d. normal random variable. For empirical purposes, in the crisis mode we assume that there is no interest rate smoothing and the policy instrument may respond to the credit spread:

$$i_t = \gamma_\pi \pi_t + \gamma_y y_t + \gamma_\omega \omega_t + c_i \varepsilon_{it}. \qquad (11)$$

Such an extended Taylor rule specification was proposed by Taylor, and analyzed by Curdia and Woodford (2010).

Since our crisis mode actually the normal times mode, it is easy to map the two modes into an MJLQ model. In particular, we assume that most of the structural parameters are constant across modes, but that the terms in the interest rate spreads and marginal utility gaps only enter in the crisis mode. Moreover, the form of the policy rule differs somewhat across modes. To be explicit, we analyze an MJLQ model of the following form:

$$\begin{aligned}
\pi_t &= \omega_f E_t \pi_{t+1} + (1-\omega_f)\pi_{t-1} + \gamma y_t + \xi_{j_t} \Omega_t + c_\pi \varepsilon_{\pi t}, \qquad (12)\\
y_t &= \beta_f E_t y_{t+1} + (1-\beta_f)\left[\beta_y y_{t-1} + (1-\beta_y) y_{t-2}\right] - \beta_r \left(i_t - E_t \pi_{t+1}\right) + \theta_{j_t} \Omega_t + \phi_{j_t} \omega_t + c_y \varepsilon_{yt}.\\
\Omega_t &= \delta E_t \Omega_{t+1} + \omega_t\\
\omega_{t+1} &= \rho_{\omega,j_{t+1}} \omega_t + c_{\omega,j_{t+1}} \varepsilon_{\omega t+1}.\\
i_t &= (1-\rho_{1,j_t} - \rho_{2,j_t})\left(\gamma_{\pi,j_t}\pi_t + \gamma_{y,j_t} y_t\right) + \gamma_{\omega,j_t}\omega_t + \rho_{1,j_t} i_{t-1} + \rho_{2,j_t} i_{t-2} + c_{i,j_t}\varepsilon_{it}. \qquad (13)
\end{aligned}$$

Here $j_t \in \{1, 2\}$ indexes the mode at date t, with mode 1 being normal times, and we assume that a transition matrix P governs the switches between modes. Thus we have $\xi_1 = \theta_1 = \phi_1 = \gamma_{\omega,1} = 0$, while $\rho_{1,1} = \rho_{2,j} = 0$. Note that we allow the dynamics of the spread ω to differ across modes both in terms of its persistence and volatility, which is key for explaining and interpreting the data. Simply put, crises are times of substantially larger volatility in interest rate spreads.

3.3. Calibration and Estimation

In this section we discuss how we fit the model to the data. We wanted to be sure to obtain estimates consistent with our interpretation of the modes, so we chose a mixture of calibration and estimation. Thus we take these estimates as suggestive for our optimal policy exercises, but make no claim to providing a full empirical analysis of the model.

We obtained all data from the St. Louis Fed FRED website. For the basic time series, we use the standard definitions: the growth of the GDP deflator is our measure of inflation, the deviation between actual GDP and the CBO estimate of potential is our measure of the output gap, and the federal funds rate is our policy interest rate. There were no significant trends overall in the data, but we do take out their means. In Figure 1 we plot these quarterly data for the period 1978:1-2011:2. We focus mostly on the Volcker-Greenspan-Bernanke era, but include some earlier data as well. The graph clearly shows the overall downward trend in inflation and nominal interest rates over this period, with the recessions of the early 1980s and the most recent period showing as large negative output gaps. For the interest rate spread, we consider two alternative indicators. The first is the gap between the yield on 3-month CDs and the federal funds rate, which is one of the spreads considered by Taylor and Williams (2008). As a somewhat broader measure of firm financing, we also consider the Option-Adjusted Spread of the BofA Merrill Lynch US Corporate A Index. For the CD spreads, we removed the mean over the whole sample. However the corporate spread data are only available from 1996 on, so for this series we subtracted the mean over the 1996-2006 period. These data are shown in Figure 2. Both series show a substantial increase in spreads starting in 2007 and peaking at the end of 2008. However the longer CD spread series also shows an earlier episode with a substantial negative spread in mid-1980. Although the spike in the corporate spread appears more dramatic, the corporate spread is more volatile overall, so the CD spread spike is roughly as much of an outlier.

Clearly we only have at most two real observations on episodes with substantial interest rate spreads, so the data won't provide much guidance in choosing among alternative specifications. In addition, it is questionable whether the large negative spreads in the 1980s were driven by similar factors as the recent large positive spreads. Certainly our interpretation of the events as financial crises does not fit with the early 1980s, when the large negative spreads were more likely the consequence of an inverted term structure than increases in liquidity or default premiums. We choose to model the interest rate spread as an AR(1) process with a switching persistence and variance, but certainly alternative specifications are plausible. This highlights another dimension of uncertainty that is not captured by our simple benchmark MJLQ model: uncertainty over the specification and evolution of the credit spreads.

In order to estimate the model, we use the methods in Svensson and Williams (2007b) to solve for an equilibrium in an MJLQ model with an arbitrary instrument rule. When we estimate the model we assume that policymakers and the public observe the current mode, although later we use these same structural parameter estimates to consider cases when the modes are unobservable. We estimate the model with Bayesian methods, finding the maximum of the posterior distribution.[9] The priors we use are discussed in Appendix A.

However, rather than simply fitting the full model to the data, in order to be sure the estimates aligned with our interpretation, we used the following approach. First, we fit the Lindé model with constant coefficients to the data for the period 1985-2006. Note that the credit spread has no interaction with the inflation and output in this mode, and thus the parameter δ is irrelevant. We deliberately cut off the beginning and end of the sample when the CD spreads were largest and most volatile, so this period represents the mode in "normal times." In addition, our model has difficulty accounting for the Volcker disinflation, which is why we chose to start only in 1985. One alternative would be to use a longer sample but to take out the trends in the data. We also estimated the model over the 1980-2006 period on

[9] We avoid saying "posterior mode" since we use "mode" in a different sense throughout the paper.

detrended data, which yielded similar results. In addition, we obtained similar results when using the corporate spread for the shorter available sample.

In our next step, we fix these estimates from the constant coefficient model as the coefficients for mode 1 (as well as the structural coefficients in mode 2) in our MJLQ model. Then we estimate the remaining parameters of the MJLQ model over the full sample from 1985-2011. As in our discussion above, we view the early 1980s episode with high interest rate spreads as arising from a separate mechanism, and so only focus on obtaining estimates of the most recent crisis. In this latter stage we are only estimating $(\xi_2, \theta_2, \phi_2, \delta, \rho_{\omega,2}, c_{\omega,2}, \gamma_{y,2}, \gamma_{\pi,2}, \gamma_{\omega,2}, c_{i,2})$ and the transition matrix P. Our estimates are given in Table 1. Our estimated transition matrix is:

$$P = \begin{bmatrix} 0.9961 & 0.0039 \\ 0.0352 & 0.9648 \end{bmatrix}.$$

Thus we see that the baseline model has a significant weight on forward looking expectations for inflation, but quite a bit less for output. The standard deviations of the shocks to inflation and the output gap are roughly equal, as is the interest rate shock in normal times. However in the crisis mode the interest rate shocks are substantially more volatile. As we'll see below, this is likely at least in part due to the fact that we do not impose the zero bound on interest rates, and thus the estimated policy rule implies negative nominal rates for the past couple of years. In the crisis mode, ξ is fairly substantial, meaning that the marginal utility gap Ω_t has a sizeable instantaneous effect on inflation, while θ is somewhat smaller. Both are positive, so Ω_t increases inflation and the output gap. The interest spread ω_t has a large negative impact on the output gap through ϕ, and spreads are substantially more volatile (and of nearly the same persistence) in the crisis mode. Finally, the crisis mode is much less persistent than the normal times mode, and the stationary distribution implied by the Markov transition matrix puts probability 0.8995 on normal times and 0.1005 on crises.

In Figure 3 we plot the estimated (filtered) probability of being in the crisis mode at each date, conditional on observations up to that date. For comparison, we also plot the CD spread once again (here scaled by 0.5 to make the scales commensurable), and for ease of interpretation we focus on the last fifteen years of data. We also plot the smoothed (two-sided) probabilities, which use the full sample to estimate the chance that the economy was in a crisis state at any given date. Here we see that these probabilities pick out exactly the crisis episode of very large magnitude spreads that we highlighted above. The filtered probabilities are rather sharp, with only small some fluctuations, but in the recent crisis there appears to be somewhat of a delay in detection. The initial run-up in CD spreads begins in mid-2007 and is interrupted by one negative observation, so the probability of a crisis mode is not clear until nearly the peak in CD spreads. Inference on the modes sharpens somewhat more when using the smoothed (two-sided) probabilities. Here we see that with the benefit of hindsight, the estimates suggest that the crisis mode began in late 2007 and ended in early 2009.

In late 2007, the filtered probability of a crisis is very low while the smoothed probability jumps up substantially. For example, in 2007:Q4-2008:Q3 the filtered probabilities of a crisis are (0.002, 0.090, 0.572), while the smoothed probabilities are (0.532, 0.687, 0.998). However this does not mean that the model has very low likelihood. Recall that process for the interest rate spread differs in normal and crisis times by having a different autocorrelation and a different variance, with the variance being especially important. Thus at each date the filtering and smoothing exercises essentially reduce to trying to determine whether a given observation is more consistent with a high or low variance mode. But even with the substantial differences in variances that we estimate (standard deviations of the interest shocks of 0.19 in normal times and 0.57 in crises), there is significant overlap in the likelihoods conditional on each mode. Thus the model initially reads the interest spread observations in late 2007 as reflecting larger shocks than the smoothed probabilities would suggest. But even these are not extreme outliers, being equivalent to observations 1-1.5 standard deviations above the predicted mean.

Overall, these results highlight that even though the probabilities of the modes appear rather sharply estimated, that there still may be uncertainty and delay in the detection of a crisis. In our initial policy analysis we will assume that all agents, both public and private, observe the current mode. But later we show how uncertainty over the current modes can change policy decisions.

4. Optimal monetary policy with financial uncertainty

4.1. Optimal policy: Observable modes (OBS)

Our MJLQ model (12) fits into the general form (1)-(2) discussed above. In particular, we have three forward-looking variables ($x_t \equiv (\pi_t, y_t, \Omega_t)'$) and consequently three Lagrange multipliers ($\Xi_{t-1} \equiv (\Xi_{\pi,t-1}, \Xi_{y,t-1}, \Xi_{\omega,t-1})'$) in the extended state space. We can write the system with seven predetermined variables: $X_t \equiv (\pi_{t-1}, y_{t-1}, y_{t-2}, i_{t-1}, \varepsilon_{\pi t}, \varepsilon_{yt}, \omega_t)'$. We use the following loss function:

$$L(X_t, x_t, i_t) = \pi_t^2 + \lambda y_t^2 + \nu(i_t - i_{t-1})^2, \tag{14}$$

which is a common central-bank loss function in empirical studies, with the final term expressing a preference for interest rate smoothing. We set the weights to $\lambda = 0.5$ and $\nu = 0.5$, and fix the discount factor in the intertemporal loss function to $\delta = 1$. We briefly discuss the role of alternative preference parameterizations below.

Then using the methods described above, we solve for the optimal policy functions

$$i_t = F_j \tilde{X}_t,$$

where now $\tilde{X}_t \equiv (\pi_{t-1}, y_{t-1}, y_{t-2}, i_{t-1}, \varepsilon_{\pi t}, \varepsilon_{yt}, \omega_t, \Xi_{\pi,t-1}, \Xi_{y,t-1}, \Xi_{\omega,t-1})'$. Thus the optimal policy consists of mode-dependent linear policy functions. It is difficult to interpret the functions directly, so we look at the implied impulse response functions.

The impulse responses of inflation, the output gap, and the interest rate to the interest rate spread are shown in Figure 4. We also plot the impulse responses under the optimal policy for the constant coefficient models which would result if the economy were to remain forever in mode 1 or mode 2. In particular, Figure 4 shows the distribution of responses from two sets of 10,0000 simulations of the MJLQ model. We initialize the Markov chain in one of the two modes and then draw simulated values of the Markov chain, plotting the median and 90% probability bands from the simulated impulse response distribution. The distribution is not apparent in the left column, as there we initialize in mode 1 which is very highly persistent, and very few of the 10,000 runs experienced a switch in the mode within the first 30 periods. The average duration of the crisis mode 2 is significantly shorter, so the right column shows the effects of some of the mode switches.

The only policy-relevant uncertainty in this model is in the response to interest rate spreads ω_t. These spreads are exogenous, and in mode 1 they do not affect inflation or the output gap. Thus in the constant-coefficient model corresponding to mode 1, there is no response of policy to the interest spread. In the constant-coefficient model corresponding to mode 2, positive interest rate spreads lead to a very sharp reduction in the output gap, and policy responds to interest rate spread shocks by sharply cutting interest rates. However as the spreads are directly observable, no other policy response is affected. The impulse responses to inflation and output gap shocks, are not shown but are the same across modes. Inflation and the output gap both jump with their own shocks, while they follow hump-shaped responses to each other's shocks. The optimal policy response is to increase interest rates in response to shocks to inflation and the output gap, with the peak response coming after three quarters.

The MJLQ optimal policies effectively average over the two constant-coefficient policies. In mode 1 of the MJLQ model there is a very small negative policy response to interest spread shocks, owing to the fact that there is a small probability in each period that the economy will switch into the crisis mode. Similarly, the response to spread shocks in mode 2 is only slightly more muted than in the corresponding constant-coefficient model, as crises are expected to be shorter lived. The impulse responses in Figure 4 show the dynamic implications of these results. The left column of panels shows the responses in normal times, where we clearly see that there is no response in the constant-coefficient case and very small responses (note the scale) in the MJLQ model. Interest rates are cut in normal times in response to an interest spread shock, but by hundredths of a basis point. By contrast, in the crisis mode interest rates are cut sharply in response to a shock, with the output gap falling and inflation increasing. We see that the median MJLQ response is nearly identical to the constant-coefficient case, but some of the mass of the distribution incorporates exits from the crisis mode, and thus corresponds to smaller responses.

4.2. Counterfactual policy simulations

In order to get a better sense of how the estimated and optimal policies may have resulted in different economic performance, we now consider some counterfactual policy experiments.

To do so, we first extract estimates of the observed Markov chain j_t and the structural shocks $(\varepsilon_{\pi t}, \varepsilon_{yt}, \varepsilon_{\omega t})$ and the policy shock ε_{it} given our estimated policy rule and structural parameters. To do so, we set the chain $j_t = 1$ if the smoothed probability (using the full sample inference) of mode 1 is greater than 0.5 and $j_t = 2$ otherwise. Then given the estimated Markov chain j_t series, we define the ε_t shocks as the residuals between the actual data and the predictions of our MJLQ model using the estimated policy rule. To consider the implications of alternative policies, we then feed the series for the Markov chain and the structural shocks through the model, zeroing out the policy shocks.

In Figure 5 we plot the simulated time series for inflation, the output gap, and the policy interest rate under the estimated monetary policy rule using the estimated shock series. For comparison, we also plot the actual data. To make the figures more interpretable, we add back in the unconditional means of the time series which we had removed for estimation. Here we see that the model tracks the data reasonably well, apart from the mid-2000s which experienced higher inflation, higher interest rates, and a higher level of the output gap than the model predicts. In general, the output gap fluctuations are more severe under the estimated policy than in the data, with the model seeming to track the fluctuations in interest rates with a lag. The model does match the decline in output and inflation over the crisis quite well, and also captures the rapid fall in interest rates. The violation of the zero lower bound is apparent over the last several quarters, as the estimated policy rule implies a fairly substantial negative interest rate.

In Figure 6 we plot similar series, but now showing the results under the optimal policy as well as those under the estimated policy rule. Here we see that the optimal policy leads to a substantial reduction in fluctuations. This is particularly true for the inflation rate, which is unsurprising since inflation fluctuations receive the largest weight in the loss function, but the cyclical fluctuations in the output gap are much more moderate as well. In the mid-1990s and again in the mid-2000s, the optimal policy calls for an earlier tightening, with interest rates beginning to increase several quarters earlier than under the estimated policy, which contributes to the lessening of inflation and output fluctuations. In the most recent crisis, the optimal policy largely follows the estimated one, with interest rates falling rapidly from mid-2008 through 2009. Under the optimal policy, this large reduction in rates leads to a massive violation of the zero lower bound on nominal rates, as the federal funds rate falls to a low of -4.36% in mid 2009. This rapid interest rate reduction under the optimal policy leads to a sharp increase in inflation, and a more moderate decline in output than under the estimated policy rule. The overall implications of the optimal policy seem to be largely to increase rates more rapidly in times of expansion, but then cut them dramatically and rapidly in crisis episodes. However the failure to incorporate the zero bound seems to be a severe constraint in taking these implications too seriously. In the next section we address one way to deal with the zero bound, and so to provide more credible policy implications.

4.3. Coping with the zero lower bound on nominal interest rates

It is difficult to directly incorporate the zero lower bound on nominal interest rates in our setting, as the bound introduces a nonlinearity which would require alternative solution methods. Eggertsson and Woodford (2003) develop one means of incorporating the zero bound and still using largely linear methods, but it is difficult to adapt their approach to our MJLQ setting. Thus rather than directly addressing the zero bound, we instead follow the approach of Woodford (2003) and incorporate an additional interest rate volatility penalty term in the loss function as a means of making the zero bound less likely to be violated. Moreover, as the zero bound is much more of a problem in crisis states, we specify that this penalty increases in the crisis mode. Thus we now use the following loss function:

$$L(X_t, x_t, i_t) = \pi_t^2 + \lambda y_t^2 + \nu(i_t - i_{t-1})^2 + \psi_{j_t} i_t^2, \tag{15}$$

where ψ_j is now the mode-dependent penalty on interest rate volatility (rather than interest smoothing). We keep the other loss function parameters the same as previously, but now set $\psi_1 = 0.7$, and $\psi_2 = 0.875$. Thus the penalty for interest rate volatility is 25% larger in the crisis state. Admittedly, giving interest rate volatility a symmetric penalty is not an entirely satisfying way to deal with the inherent asymmetries that zero bound introduces. Nonetheless, this penalty does ensure that the bound is satisfied in the sample we consider.

The optimal policies with the interest rate penalties are largely similar to our previous results. However because the loss function now varies across modes, policy responses to all variables change with the mode, if only slightly. Thus the switching penalty slightly muddies our previous result that only the response to interest rate spreads changed in crises. The increased interest rate penalty in crisis times means that the responses to all variables except the interest rate spread are more muted in mode 2 than in mode 1. In the MJLQ model policymakers also now anticipate that the penalty for interest rate volatility will switch with the mode, which affects (at least slightly) their responses to all variables. However it remains the case that the mode-dependent MJLQ responses are very similar to the corresponding constant-coefficient responses, especially in the highly persistent normal times mode. We now take this specification with a switching interest rate volatility penalty as our baseline.

In Figure 7 we show the counterfactual time series for inflation, the output gap, and the nominal interest rate for the two optimal policies: the previous case with only an interest smoothing term in the loss function and the current case with a switching interest rate volatility penalty as well. We also show the results under the estimated policy rule and the actual data for comparison. Overall, adding the interest rate volatility penalty dampens the fluctuations in the nominal interest rate, just as one would expect. However the effects on inflation are quite modest for most of the sample. The interest rate path for the optimal policy with the interest volatility penalty actually seems to match the actual data fairly well, especially over the period from about 2000-2009. However inflation increases in the crisis under the optimal policies, and the output gap declines are more moderate than in the data. In response, the optimal policies call for increases in interest rates over the last couple of years of the sample. Even though the policy instrument paths are similar over that time span, the expectations channel seems to play an important role in driving the differences in outcomes. The optimal policy with an interest volatility penalty is significantly more accommodative of higher inflation, so the reductions in interest rates beginning in 2007 lead to expectations of higher inflation going forward. Thus the optimal policy is able to generate a substantial inflation during the crisis, which leads to a much more modest decline in output than was actually experienced.

Overall, we have seen that the switches in modes from normal times to crises has a large effect on policy, but the uncertainty about the future switches has relatively little effect. In crises, it is optimal to cut interest rates substantially in response to increases in the interest rate spread. However the size of this response is nearly the same in our MJLQ model as in the corresponding constant coefficient model. In addition, the possibility that the economy may enter a crisis means that even in normal times policy should respond to interest rate spreads. But again, this effect is fairly negligible.

These results seem to rely on the exogeneity of the interest rate spreads, as well as the rarity of crises. In regard to the first point, policy cannot affect spreads in our model, so responding to interest rate spreads in normal times has no effect on the severity of crises. If policy could affect spreads, then there may be more of a motive to react before a crisis would appear, as stabilizing interest spreads may make crises less severe. On the second point, note that by responding to spreads in normal times policymakers are effectively trading off current performance for future performance. The greater the chance of transiting into a crisis, the larger the weight that the uncertain future would receive in this tradeoff. As the normal times mode is very highly persistent in our estimates, there is little reason to sacrifice much current performance for the chance of moderating future crises. Once a crisis starts, policymakers are able to cut interest rates sufficiently to help stabilize the economy, even when there are costs to interest rate volatility.

4.4. Optimal policy: Unobservable modes (NL and AOP)

We now turn to the case where there is uncertainty about the current state of the economy, and agents must learn whether a crisis has begun. This could potentially increase the precautionary motive for policy, increasing the response to financial factors in normal times. When we focused on the observable case above, we assumed that the shocks $(\varepsilon_{\pi t}, \varepsilon_{yt}, \varepsilon_{\omega t})$ were observable and policy could respond directly to them. However, to focus on the role of learning, we now assume that the shocks are unobservable. If they were observable, then agents would be able to infer the mode from their observations of the forward-looking variables and the interest rate spread.

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

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

where now $\tilde{X}_t \equiv (\pi_{t-1}, y_{t-1}, y_{t-2}, i_{t-1}, \omega_t, \Xi_{\pi,t-1}, \Xi_{y,t-1}, \Xi_{\omega,t-1})'$. In addition, we must track the estimated mode probabilities $(p_{t|t} \equiv (p_{1t|t}, p_{2t|t})'$, of which we only need keep track of $p_{1t|t})$. Thus, the value and policy functions are nine dimensional. Computational constraints thus prohibit us from solving for the full value functions in the AOP case, and prevent us from considering the BOP case at all. However we can still fully solve for the NL case and implement the AOP case recursively. Additionally, Monte Carlo simulation allows us to evaluate losses under NL and AOP as well.

In Figure 8, we plot the impulse responses initialized in mode 2. The figure plots the results of 10,0000 simulations (of the Markov chain) of the impulse responses to shocks when the modes are unobservable. We initialize the Markov chain in the crisis mode 2, and set the initial beliefs $p_{0|0}$ at the stationary distribution of the Markov chain P. Since there is only a single shock to learn from, the AOP and NL impulse responses were all essentially identical, so we only plot the AOP responses. For comparison, we also plot the responses with observable modes, as in Figure **??** above. Overall, we see that the responses of most of the variables are more sluggish and muted when the current mode is unobservable. This is perhaps most clear in the response of the interest rate to a credit spread shock. Rather than cutting interest rates upon impact of the shock, as happens in the observable case, there is initially essentially no response because the beliefs put very high probability of being in normal times. Only after the shock works its way through the economy, starting with the decline in output and increase in inflation on impact, does the interest rate respond. Note also that the response of inflation to the credit spread shock when the modes are unobservable is only about half as large as in the observable case. This suggests the importance of uncertainty for private sector behavior as well as policy decisions.

In Figure 9 we show the counterfactual time series for inflation, the output gap, and the nominal interest rate when the modes are observable and the two unobservable cases of no learning and adaptive optimal policy. Overall, the fluctuations in variables are larger in the observable case, which is particularly noticeable for the output gap and interest rates. As the economy was in normal times throughout most of the sample, there is essentially no difference between the NL and AOP results until late 2007 when the economy switched into the crisis mode. Interestingly, the no learning case seemed to perform best in that episode, as the increase in inflation was substantially smaller and the fall in output slightly lower than under AOP or in the observable case. As the beliefs of both private agents and the central bank remain constant at the stationary distribution in the NL case, there is much less responsiveness of all variables to the crisis. In the AOP case, the switch gets discovered relatively quickly, and inflation and the output gap more closely follow the observable counterpart. Surprisingly however, the interest rate responds the least in AOP case. In summary, uncertainty about the current regime of the economy has relatively minor effects. Agents are able to detect crises rather quickly once they observe large increases in interest spreads, so the AOP outcomes are close to the observable case.

5. Conclusion

This paper has illustrated how to formulate and analyze monetary policy with uncertainty about the impact of the financial sector on the broader economy. We have found that uncertainty about financial crises causes substantial changes in optimal monetary policies, but such changes are mostly due to the crises and not the uncertainty. In our estimated model, crises are infrequent, exogenous events and so policy in normal times is affected relatively little by the possibility of crises. In addition, even if crises are not directly observable, they are relatively easy to detect, so uncertainty and learning about the state of the economy play a relatively minor role. We find that policy should indeed by tailored to crises, but that such considerations are largely independent of how policy should be conducted in normal times.

Of course, these conclusions are specific to the particular model that we analyze. In addition, even in the context of this model, the dimensions of uncertainty we consider are rather limited. Policymakers and private agents know the form and severity of crises, and they know the expected frequency and durations of crisis episodes. Thus we certainly understate the degree of uncertainty that policymakers face. We have carried out some preliminary exercises analyzing uncertainty about the duration of crises, which we implemented by having separate crisis modes of different persistence, and uncertainty about the severity of crises, implemented by having separate crisis modes with different values for the key parameters governing the financial frictions. While these increased the impact of uncertainty, the effects were rather minor. More important is likely to be a broader role for financial frictions.

While the version of the model of Curdia and Woodford (2009b) that we use is a simple staring point, it incorporates financial frictions in a limited way. Most prominently, we have focused on a version of the model where the key credit spread is exogenous. Curdia and

Woodford develop a more general version in which this spread evolves endogenously and is dependent on the level of private borrowing, which in turn depends on interest rates. The role of monetary policy in mitigating crises may be larger when policymakers have some control over interest spreads. More broadly, the model abstracts from investment, which is a key channel in the financial accelerator model of Bernanke, Gertler, and Gilchrist (1999). In their model financial frictions entail an important role for business balance sheets, which in turn makes aggregate net worth a key state variable. More recent work by Christiano, Motto, and Rostagno (2009), includes many of the real and nominal frictions studied by Smets and Wouters (2003) and Christiano, Eichenbaum, and Evans (2005), along with the financial frictions of Bernanke, Gertler, and Glichrist (1999) embedded into an explicit banking sector. The financial frictions play a more prominent role in the transmission mechanism in these models, and so the policy reactions to financial variables may be even more crucial in such settings.

There are many related issues which can also be addressed using our approach. For example, while we focused on uncertainty about the impact of financial frictions, there is also uncertainty about the type of frictions which best describes the economy. We could embed some of the models just discussed as alternative possible modes. The types of frictions – real, nominal, and financial – and their interactions vary substantially across these models, and thus policy implications may differ substantially as well. Finally, one important aspect of the current crisis has been that policymakers have engaged in "unconventional" policies, including purchases of a broad range assets and direct lending to the private sector. The models of Curdia and Woodford (2009a) and Gertler and Karadi (2009) allow for such additional channels of policy response. By embedding such models in our setting we can analyze how these unconventional instruments should be used in an uncertain environment, and how they would interact with the more conventional policies.

In all of these cases, the MJLQ approach provides a simple and flexible way of structuring and analyzing optimal policy under uncertainty. By appropriately specifying the structure, the MJLQ framework can provide guidance to policymakers on how to deal with the broad forms of uncertainty they face.

References

Bernanke, B. S., Gertler, M., 1989. Agency Costs, Net Worth, and Business Fluctuations. American Economic Review 79, 14–31.

Bernanke, B. S., Gertler, M., Gilchrist, S., 1999. The Financial Accelerator in a Quantitative Business Cycle Framework. In: Taylor, J., Woodford, M. (Eds.), Handbook of Macroeconomics. Vol. 1C. North-Holland, Amsterdam, pp. 1341–1390.

Blake, A. P., Zampolli, F., 2006. Time Consistent Policy in Markov Switching Models, working Paper No. 298, Bank of England.

Christiano, L., Eichenbaum, M., Evans, C., 2005. Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. Journal of Political Economy 113, 1–45.

Christiano, L. J., Motto, R., Rostagno, M., 2003. The Great Depression and the Friedman-Schwartz Hypothesis. Journal of Money, Credit and Banking 35, 1119–1198.

Christiano, L. J., Motto, R., Rostagno, M., 2009. Financial Factors in Economic Fluctuations, working paper, Northwestern University.

Cogley, T., Colacito, R., Sargent, T. J., 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.

Curdia, V., Woodford, M., 2009a. Conventional and Unconventional Monetary Policy, working paper, Columbia University.

Curdia, V., Woodford, M., 2009b. Credit Frictions and Optimal Monetary Policy, working paper, Columbia University.

Curdia, V., Woodford, M., 2010. Credit Frictions and Monetary Policy. Journal of Money, Credit and Banking 42, s3–35.

Eggertsson, G., Woodford, M., 2003. The Zero Bound on Interest Rates and Optimal Monetary Policy. Brookings Papers on Economic Activity 1, 139–211.

Ellison, M., 2006. The Learning Cost of Interest Rate Reversals. Journal of Monetary Economics 53, 1895–1907.

Ellison, M., Valla, N., 2001. Learning, Uncertainty and Central Bank Activism in an Economy with Strategic Interactions. Journal of Monetary Economics 48, 153–171.

Evans, G. W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press.

Gertler, M., Gilchrist, S., Natalucci, F., 2007. External Constraints on Monetary Policy and the Financial Accelerator. Journal of Money, Credit and Banking 39, 295–330.

Gertler, M., Karadi, P., 2009. A Model of Unconventional Monetary Policy, working paper, NYU.

Giannoni, M., 2002. Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model. Macroeconomic Dynamics 6, 111–144.

Kiyotaki, N., Moore, J. H., 1997. Credit Cycles. Journal of Political Economy 105, 211–248.

Levin, A. T., Onatski, A., Williams, J. C., Williams, N., 2006. Monetary policy under uncertainty in micro-founded macroeconometric models. In: Gertler, M., Rogoff, K. (Eds.), NBER Macroeconomics Annual 2005. MIT Press for the NBER, Cambridge, pp. 229–287.

Levin, A. T., Wieland, V., John C. Williams, 2003. The performance of forecast-based monetary policy rules under model uncertainty. American Economic Review 93, 622–645.

Lindé, J., 2005. Estimating New-Keynesian Phillips Curves: A Full Information Maximum Likelihood Approach. Journal of Monetary Economics 52, 1135–1149.

Marcet, A., Marimon, R., 1998. Recursive Contracts, working paper, Pompeu Fabra.

Onatski, A., Stock, J. H., 2002. Robust monetary policy under model uncertainty in a small model of the US economy. Macroeconomic Dynamics 6, 85–110.

Smets, F., Wouters, R., 2003. An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. Journal of the European Economic Association 1, 1123 – 1175.

Svensson, L. E., Williams, N., 2007a. Bayesian and Adaptive Optimal Policy Under Model Uncertainty, working paper, University of Wisconsin.

Svensson, L. E., Williams, N., 2007b. Monetary Policy with Model Uncertainty: Distribution Forecast Targeting, working paper, University of Wisconsin.

Taylor, J. B., Williams, J. C., 2008. Further Results on a Black Swan in the Money Market, sIEPR Discussion Paper No. 07-046.

Tesfaselassie, M. F., Schaling, E., Eijffinger, S. C., 2006. Learning about the Term Structure and Optimal Rules for Inflation Targeting, cEPR Discussion paper No. 5896.

Woodford, M. D., 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press, Princeton.

Table 1: Estimates of the benchmark MJLQ model.

Parameter	Mode 1	Mode 2
ω_f	0.5827	0.5827
γ	0.0137	0.0137
ξ	0	0.6468
β_f	0.2449	0.2449
β_y	0.9533	0.9533
β_r	0.0614	0.0614
θ	0	0.2802
ϕ	0	-1.6152
δ	0.2932	0.2932
ρ_w	0.4930	0.4959
ρ_1	0.8715	0
ρ_2	0.0044	0
γ_π	1.6897	0.8039
γ_y	1.1033	0.4320
γ_ω	0	-0.6819
c_π	0.4646	0.4646
c_y	0.4349	0.4349
c_ω	0.1889	0.5688
c_i	0.4484	1.2034

Figure 1: The key economic data on inflation, the output gap, and interest rates, 1978:1-2011:2

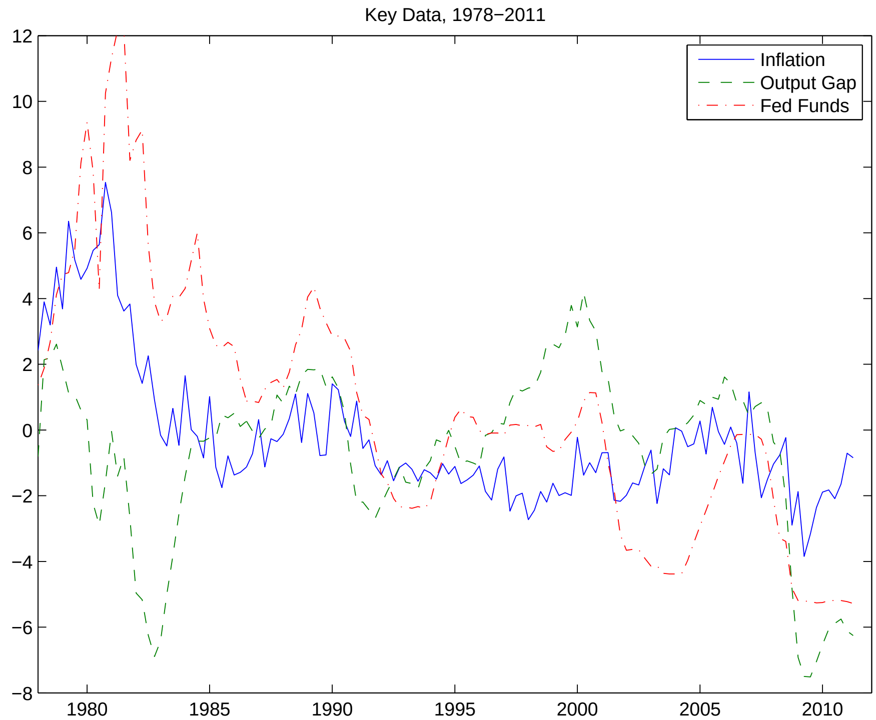

Figure 2: Two interest rate spread time series, 1978:1-2011:2

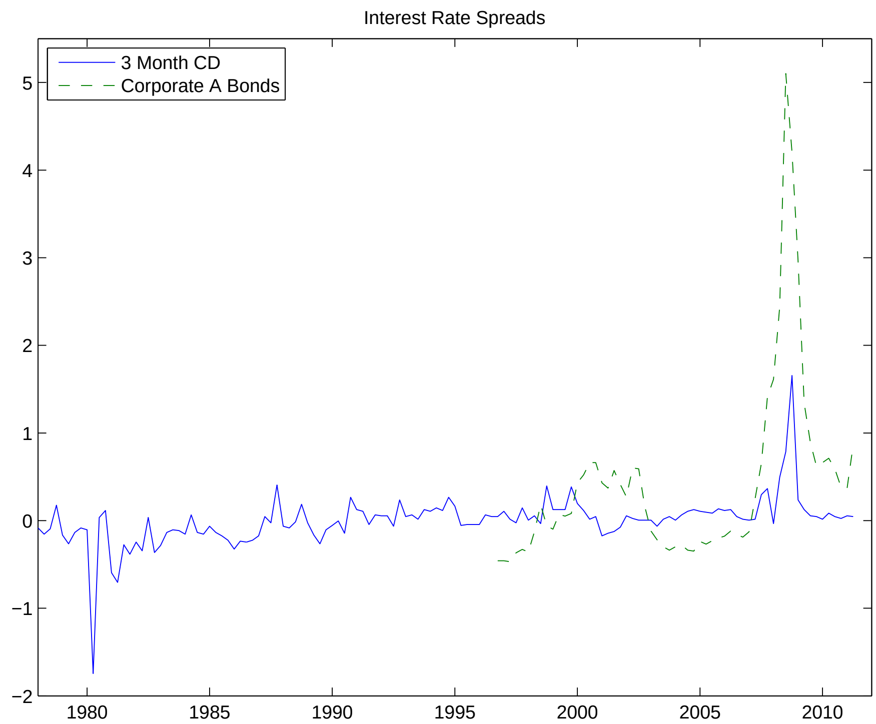

Figure 3: Probability of being in a crisis mode and 0.5*CD spread, 2001-2011:2

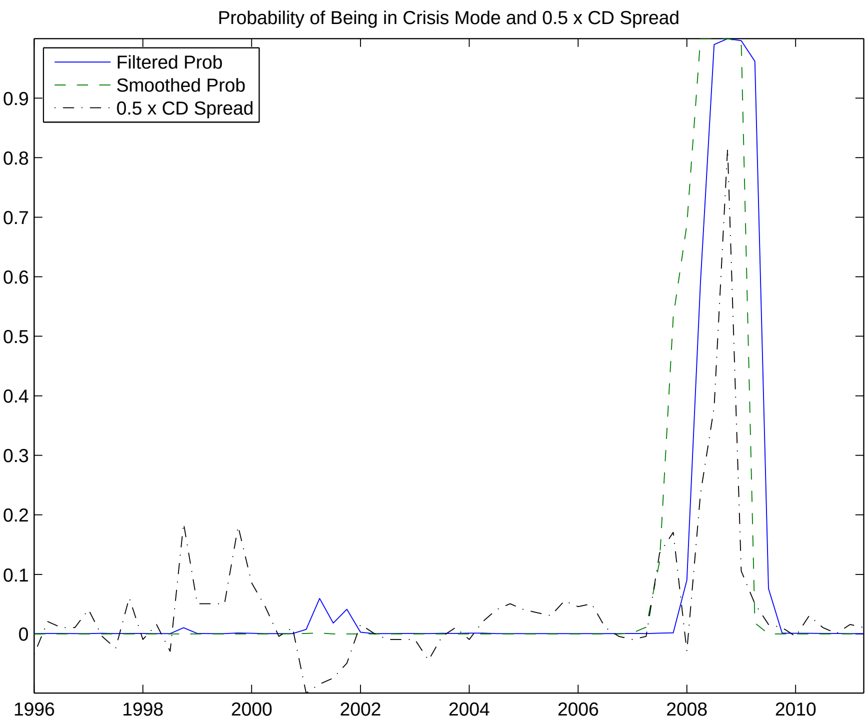

Figure 4: Impulse response of selected variables to an interest spread shock, starting in mode 1 (left column) or mode 2 (right column).

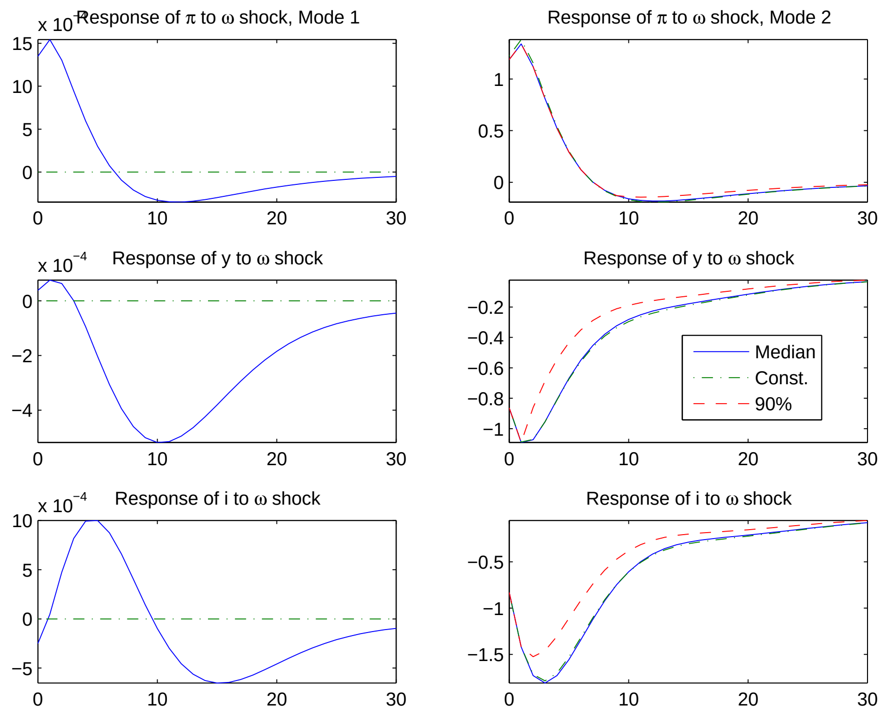

Figure 5: Simulation of the economy under the estimated policy rule (solid line) using the estimated shocks, along with actual data (dashed line).

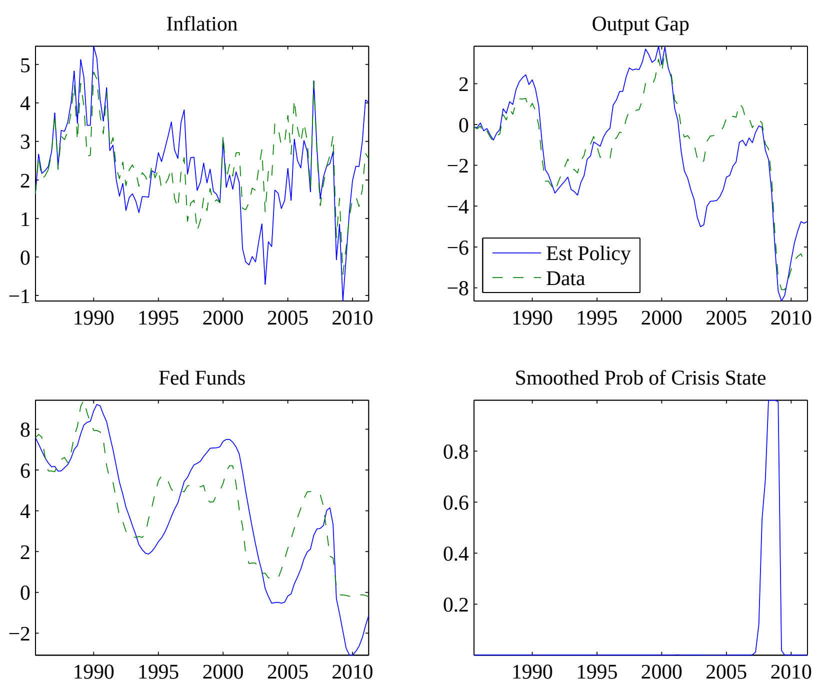

Figure 6: Simulation of the economy under the optimal policy rule (solid line) and the estimated policy rule (dashed line) using the estimated shocks.

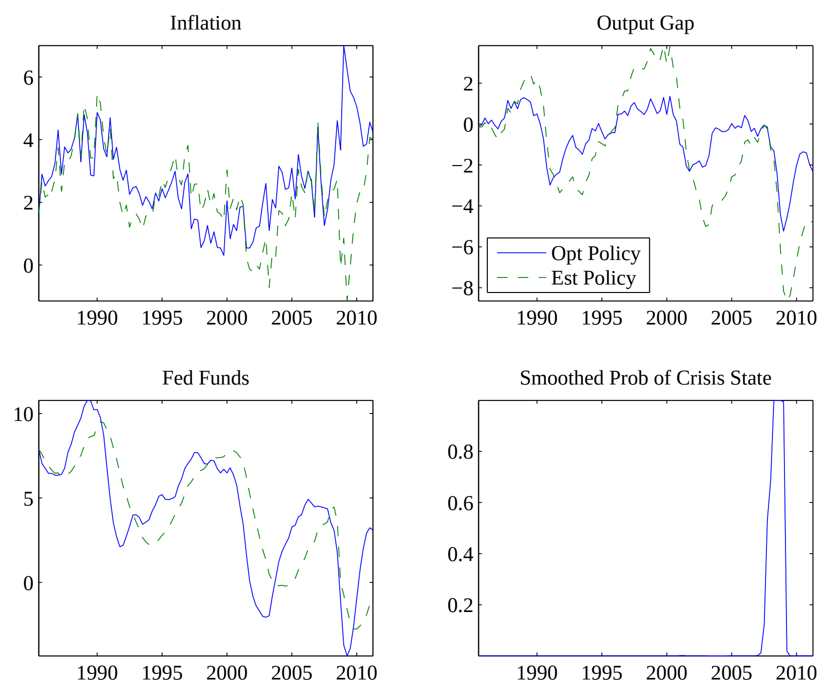

Figure 7: Simulation of the economy under the optimal policy rule with an interest volatility penalty (solid line), the optimal policy with an interest smoothing penalty (dot-dash), and the estimated policy rule (dash) using the estimated shocks, along with the actual data (dotted).

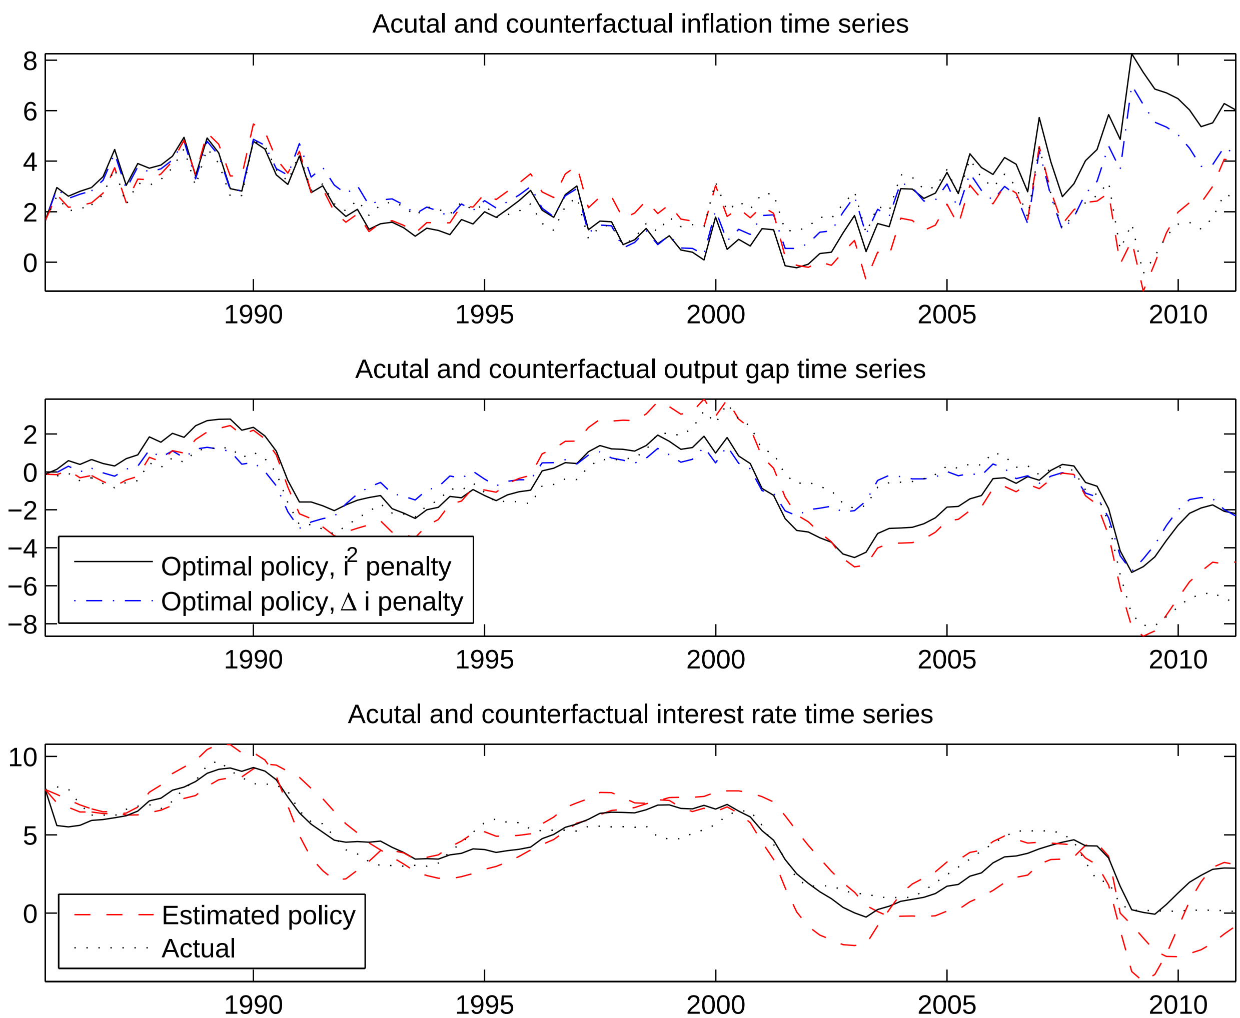

Figure 8: Impulse response of selected variables to inflation, output gap, and interest spread shocks. Impulses when the modes are unobservable (solid line) and 90% probability bands (dash), along with observable modes (dot dash). Simulations are initialized in mode 2, and beliefs are initialized at the stationary distribution.

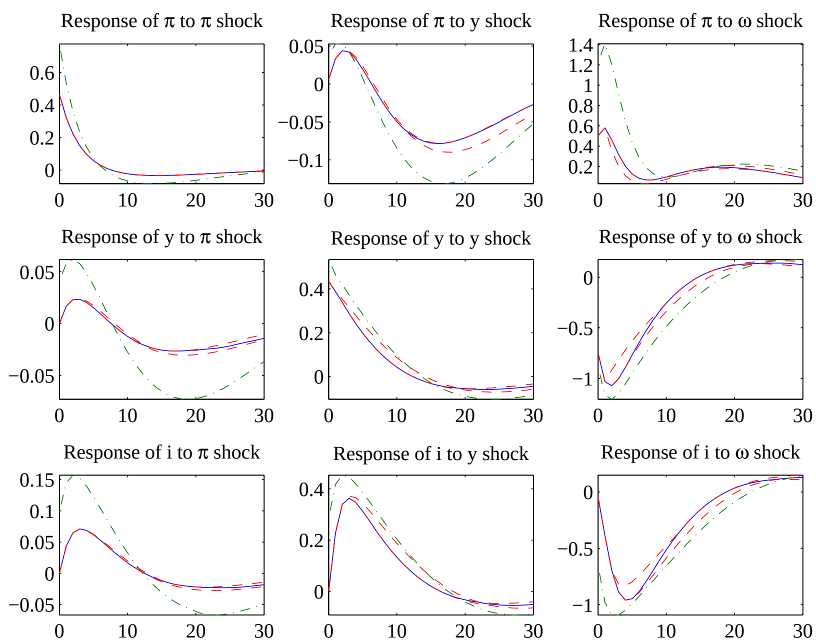

Figure 9: Simulation of the economy under the optimal policies when the modes are observable (OBS, dot-dash) as well as when they are unobservable but agents do not learn (NL, solid) or when they update beliefs and use the adaptive optimal policies (AOP, dash).