Woodford's Approach to Robust Policy Analysis in a Linear-Quadratic Framework*

Hyosung Kwon† Jianjun Miao‡

April 2013

Abstract

This paper extends Woodford's (2010) approach to the robustly optimal monetary policy to a general linear quadratic framework. We provide algorithms to solve for a time-invariant linear robustly optimal policy in a timeless perspective and for a time-invariant linear Markov perfect equilibrium under discretion. We apply our methods to a New Keynesian model of monetary policy with persistent cost-push shocks and inflation persistence. We find that the robustly optimal commitment inflation is less responsive to a cost-push shock when the shock is more persistent, and that the robustly optimal discretionary policy is more responsive to lagged inflation when inflation is more persistent.

JEL Classification: D81, D84, E52

Keywords: robustness, ambiguity, commitment, discretion, policy analysis

*We would like to thank Mike Woodford and Simon Gilchrist for helpful comments. We have benefitted from comments by the seminar participants at Boston University. First version: February 2012.

†Department of Economics, Boston University.

‡Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215. Tel.: 617-353-6675. Email: miaoj@bu.edu. Homepage: http://people.bu.edu/miaoj.

1. Introduction

The traditional policy analysis has been typically conducted under the rational expectations hypothesis. While we have learned many lessons from this hypothesis, there are several good reasons for us to think about departures from it. First, the Ellsberg (1961) paradox and related experimental evidence demonstrate that there is a distinction between risk and ambiguity. Risk refers to the situation where there is a known probability distribution over the state of the world, while ambiguity refers to the situation where the information is too vague to be adequately summarized by a single probability distribution. As a result, a decision maker may have multiple priors in mind (Gilboa and Schmeidler (1989)). Second, as Anderson, Hansen and Sargent (2003) and Hansen and Sargent (2001, 2008) point out, economic agents view economic models as an approximation to the true model. They believe that economic data come from an unknown member of a set of unspecified models near the approximating model. Concern about model misspecification induces a decision maker to want robust decision rules that work over that set of nearby models.[1]

The above arguments are especially relevant for policy analysis in which uncertainty and expectations play an important role. Either the policymaker or the private agents may not have complete confidence in the likelihood of the state of the world and hence face model ambiguity. Several possibilities may arise.[2] First, the policymaker does not trust its approximating model, but the private agents do. The policymaker forms expectations using a model in a set of nearby models surrounding its approximating model. Hansen and Sargent (2003 and 2008, Chapter 16) study this case. Woodford (2010) studies a second case in which the policymaker trusts its approximating model, but it has ambiguity about the private agents' beliefs about the model. The policymaker thinks that the private agents may form expectations using any model in a set of nearby models surrounding the policymaker's approximating model. The Hansen-Sargent approach seems well understood and several computation algorithms in the linear-quadratic framework have been proposed in the literature (e.g., Hansen and Sargent (2008), Giordani and Soderlind (2008), Dennis (2008), Leitemo, and Soderstrom (2009)). However, little is known about how to conduct robust policy analysis using the Woodford approach in a general linear-quadratic framework.

This paper fills this gap by proposing algorithms for solving robustly optimal policy under both timeless perspective commitment and discretion. Our algorithms borrow insights from

[1] There is a growing literature on the applications of robustness and ambiguity to finance and macroeconomics, e.g., Epstein and Wang (1994); Hansen (2007); and Ju and Miao (2012); among others.

[2] See Hansen and Sargent (2012) for a discussion of different types of ambiguity.

Woodford (2010), who derives analytical solutions for a basic New Keynesian model. We focus on conditionally linear policy rules under timeless perspective commitment and linear Markov perfect equilibria (MPE) under discretion. We show that robustly optimal conditionally linear policy and a linear Markov perfect equilibrium can exist given certain conditions. Our algorithms then reduce to solving for linear decision rules using tools from the literature on solving linear rational expectations models and linear-quadratic control problems.

We apply our algorithms to a New Keynesian model with a persistent cost-push shock and inflation persistence. The Woodford (2010) model is a special case of our model. We confirm Woodford's (2010) major findings for the case of purely temporary shocks that (i) the robustly optimal inflation is less responsive to the current cost-push shock, and (ii) the robustly optimal inflation is more history-dependent compared to the case of rational expectations. We also find some new comparative statics results: (i) When the cost-push shock becomes more persistent, robustly optimal inflation is less responsive to it, but the robustly optimal output gap is more responsive. By contrast, optimal inflation under rational expectations is more responsive to a cost-push shock when it is more persistent. (ii) In addition, the robust discretionary policy becomes more costly and the range of parameters for the existence of a linear MPE shrinks. (iii) When the steady state is more inefficient, the robustly optimal inflation is less responsive to the cost-push shock.

As Woodford (2010) points out, the central bank is not willing to let inflation increase in response to a positive cost-push shock when it is faced with a concern for robustness. The central bank fears that a large shock in inflation might affect the inflation expectation of private agents unfavorably to the central bank and thus the output-inflation trade-off might worsen if it allows inflation to increase. Naturally, the central bank's concern for the unfavorable change in the private agents' forecast becomes larger when the cost-push shock is more persistent. As a result, the central bank is more conservative in setting the optimal inflation rate. If the shock is as highly persistent as a unit root process, the central bank commits to an inflation rate which does not depend on the current shock.

In our model, the New Keynesian Phillips curve contains lagged inflation. The lagged inflation rate is an endogenous state variable and hence the method in Woodford (2010) cannot be applied. Using our algorithms, we find the following new results: (i) When the degree of inflation inertia is larger, the sensitivities to the contemporaneous cost-push shock of both the robustly optimal commitment and discretionary policies are closer to those under rational expectations. In this case, the impact of the belief distortion from the private sector is smaller as a smaller weight is attached to the expected inflation. (ii) The robustly optimal discretionary

policy is more sensitive to lagged inflation than the optimal discretionary policy under rational expectations. This sensitivity increases as the volatility of the cost-push shock increases. This is in contrast to the case of rational expectations under which the optimal discretionary inflation responds to lagged inflation by the same degree regardless of the size of the shock volatility. The intuition is that, without commitment, the inflation bias is larger with concerns for robustness than in the case of rational expectations. Under the worst-case beliefs, inflation expectations are biased upward, and hence a discretionary central bank has a greater incentive to respond to both the current shock and lagged inflation.

The remainder of the paper proceeds as follows. Section 2 presents the general framework. Section 3 presents an algorithm for solving robustly optimal policy in a timeless perspective. Section 4 provides an algorithm for solving robustly optimal discretionary policy. Section 5 analyzes a monetary policy example. Section 6 concludes. Technical details are relegated to appendices.

2. A Linear-Quadratic Framework

2.1. Uncertainty and Beliefs

Uncertainty is generated by a stochastic process of shocks $\{\varepsilon_t\}_{t=0}^{\infty}$ where ε_t is an $n_\varepsilon \times 1$ vector of independently and identically distributed standard normal random variable. Let $\varepsilon^t = \{\varepsilon_0, \varepsilon_1, ..., \varepsilon_t\}$. At date t, both the policymaker and the private sector have common information generated by ε^t and some initial state x_0. They may not have rational expectations in that their subjective beliefs may not coincide with the objective probability distribution governing exogenous shocks $\{\varepsilon_t\}_{t=0}^{\infty}$. One reason is that economic agents view their model as an approximation and thus may be concerned about model misspecification.

Model misspecification is described by a perturbation to the distribution of shocks. Let $p(\varepsilon)$ denote the standard normal density of ε_t. Let P and P_t denote the induced distribution over the full state space and the induced joint distribution of ε^t, respectively. Assume that a distorted distribution is absolutely continuous with respect to the reference distribution P. We can then represent belief distortions by Radon-Nikodym derivatives.

Let $\hat{p}\left(\varepsilon|\varepsilon^t, x_0\right)$ denote an alternative one-step-ahead density for ε_{t+1} conditioned on date t information. Form the likelihood ratio or the Radon-Nikodym derivative for one-step-ahead distributions:

$$m_{t+1} = \frac{\hat{p}\left(\varepsilon|\varepsilon^t, x_0\right)}{p(\varepsilon)}.$$

It satisfies the property

$$E_t\left[m_{t+1}\right]=1, \tag{1}$$

where E_t denotes the conditional expectation operator for the reference distribution P given date t information. Inspired by Hansen and Sargent (2001, 2008), Woodford (2010) uses conditional relative entropy to measure the discrepancy between the distorted distribution and the reference distribution. He first constructs the conditional relative entropy of a one-step-ahead distribution given date t information as $E_t\left[m_{t+1}\ln m_{t+1}\right]$. He then defines the expected discounted entropy conditioned on date zero information as

$$E_0\sum_{t=0}^{\infty}\beta^t\left[E_t\left(m_{t+1}\ln m_{t+1}\right)\right]=E_0\sum_{j=0}^{\infty}\beta^t m_{t+1}\ln m_{t+1}. \tag{2}$$

Model ambiguity is described by a set of one-step-ahead densities $\{m_t\}_{t=1}^{\infty}$ satisfying the constraint:

$$E_0\sum_{t=0}^{\infty}\beta^t m_{t+1}\ln m_{t+1}\leq\eta_0 \tag{3}$$

for some $\eta_0>0$.

2.2. Robustly Optimal Policy Problem

Suppose that the equilibrium system from the private sector can be summarized by the following form:

$$\begin{bmatrix} I & 0 \\ D_{21} & D_{22} \end{bmatrix}\begin{bmatrix} x_{t+1} \\ \hat{E}_t y_{t+1} \end{bmatrix}=\hat{A}\begin{bmatrix} x_t \\ y_t \end{bmatrix}+\hat{B}u_t+\begin{bmatrix} \widehat{C}_x \\ 0 \end{bmatrix}\varepsilon_{t+1}, \tag{4}$$

where x_0 is exogenously given and $\hat{E}_t$ denotes the conditional expectation operator given date t information based on the common beliefs of the private sector. The private sector's beliefs may not coincide with the "objective" probability distribution for $\{\varepsilon_t\}$, the reference distribution P. Here x_t is an $n_x\times 1$ vector of predetermined variables in the sense defined in Klein (2000), y_t is an $n_y\times 1$ vector of non-predetermined or forward-looking variables, and u_t is an $n_u\times 1$ vector of instrument or control variables chosen by the policymaker. We typically use x_t to represent the state of the economy, which may include productivity shocks, preference shocks, or capital stock. Note that x_t may include a component of unity in order to handle constants. The vector y_t represents endogenous variables such as consumption, inflation rate, and output.

Examples of instruments u_t include interest rates and money growth rates. The equation for x_t is backward-looking and represents the law of motion of state variables. The equation for y_t is forward-looking and typically represents the first-order conditions from intertemporal optimization, such as Euler equations.

All matrices in (4) are conformable. For simplicity, we suppose that the matrix on the left side of equation (4) is invertible so that we can multiply both sides of this equation by its inverse to obtain the system:[3]

$$\begin{bmatrix} x_{t+1} \\ \hat{E}_t y_{t+1} \end{bmatrix} = A \begin{bmatrix} x_t \\ y_t \end{bmatrix} + Bu_t + C\varepsilon_{t+1}, \tag{5}$$

where we have used the shorthand notation,

$$A \equiv \begin{bmatrix} A_{xx} & A_{xy} \\ A_{yx} & A_{yy} \end{bmatrix}, \quad B \equiv \begin{bmatrix} B_x \\ B_y \end{bmatrix}, \quad C \equiv \begin{bmatrix} C_x \\ 0 \end{bmatrix},$$

and where we have partitioned matrices conformably.

The policymaker has the period loss function,

$$L(x_t, y_t, u_t) = \frac{1}{2}\left[x_t', y_t'\right] \begin{bmatrix} Q_{xx}, & Q_{xy} \\ Q_{xy}' & Q_{yy} \end{bmatrix} \begin{bmatrix} x_t \\ y_t \end{bmatrix} + \frac{1}{2} u_t' R u_t + \left[x_t', y_t'\right] \begin{bmatrix} S_x \\ S_y \end{bmatrix} u_t,$$

where we assume the matrix

$$\begin{bmatrix} Q_{xx} & Q_{xy} & S_x \\ Q_{xy}' & Q_{yy} & S_y \\ S_x' & S_y' & R \end{bmatrix},$$

is symmetric and positive semidefinite.

If both the private sector and the policymaker have rational expectations, then they have common beliefs which coincide with P, the probability distribution governing exogenous shocks $\{\varepsilon_t\}$. In this case, the optimal policy problem with commitment is given by

$$\max_{\{x_t, y_t, u_t\}} \; -E_0 \sum_{t=0}^{\infty} \beta^t L(x_t, y_t, u_t), \tag{6}$$

subject to (5) in which the conditional expectation operator $\hat{E}_t$ is equal to E_t, the conditional expectation operator with respect to P.

[3] The singular case can be handled by the QZ decomposition method, e.g., Klein (2000) and Sims (2002).

Following Woodford (2010), we suppose that the policymaker has a single model of the exogenous processes $\{\varepsilon_t\}$ and thus no ambiguity along this dimension. Nevertheless, the policymaker faces ambiguity because it knows only that the private sector's model is within a set of probability models surrounding its own model. The policymaker evaluates the private sector's forward-looking equation using a worst-case model and solves the following problem:

$$\max_{\{x_t,y_t,u_t\}} \min_{\{m_{t+1}\}} \quad -E_0 \sum_{t=0}^{\infty} \beta^t L(x_t, y_t, u_t) + \theta E_0 \sum_{t=0}^{\infty} \beta^t m_{t+1} \ln m_{t+1}, \tag{7}$$

subject to (1) and

$$\begin{bmatrix} x_{t+1} \\ E_t[m_{t+1} y_{t+1}] \end{bmatrix} = A \begin{bmatrix} x_t \\ y_t \end{bmatrix} + Bu_t + C\varepsilon_{t+1}. \tag{8}$$

Here the parameter $\theta > 0$ penalizes one-step-ahead densities $\{m_t\}$ with large relative entropies defined in (2). It may be regarded as the Lagrange multiplier for the constraint (3). Following Hansen and Sargent (2001, 2008), instead of solving for the constraint problem subject to (3), we treat θ as a parameter, which measures the policymaker's degree of concern for possible departures from rational expectations, with a small value of θ implying a great degree of concern for robustness, while a large value of θ implies that only modest departures from rational expectations are considered plausible. When $\theta \to \infty$, the rational expectations analysis is obtained as a limiting case.

3. Commitment in a Timeless Perspective

Instead of solving for an optimal commitment policy,[4] Woodford (2010) uses a timeless perspective to derive a closed-form solution to a robustly optimal monetary policy problem in a simple New Keynesian model. We now extend his idea to our general formulation. We replace E_0 in (7) with E_{-1}, the conditional expectation operator given the economy's state at date -1, i.e., before the realization of the period zero state x_0 or implied ε_0. Instead of supposing that the policymaker chooses a sequence of (possibly time-varying) $\{y_t\}$ for all $t \geq 0$, we consider only the problem of choosing an optimal sequence of commitments $\{y_t\}$ for periods $t \geq 1$, taking as given a commitment y_0 that the policymaker must fulfill.

Suppose that the initial commitment takes a linear form:

$$y_0 = \Phi_{-1} + \Gamma_{-1}\varepsilon_0,$$

[4] Kwon and Miao (2012) study this type of policy in models with three types of ambiguity.

for some coefficients (Φ_{-1}, Γ_{-1}). Suppose that $\{y_t\}$ chosen for periods $t \geq 1$ may depend on (Φ_{-1}, Γ_{-1}) and shocks from periods zero through t. We focus on conditionally linear rules of the following form:

$$y_{t+1} = \Phi_t + \Gamma_t \varepsilon_{t+1}, \ t \geq 0, \tag{9}$$

where Φ_t is stochastic and measurable with respect to date t information and Γ_t is deterministic. For any initial commitment (Φ_{-1}, Γ_{-1}), the policymaker chooses $(\Phi_t, \Gamma_t)_{t\geq 0}$, $\{x_t, m_t\}_{t\geq 1}$, and $\{u_t\}_{t\geq 0}$ to solve problem (7).

To define optimal policy from a timeless perspective, suppose that Φ_{-1} is drawn from some distribution ρ. The initial commitment (Φ_{-1}, Γ_{-1}) is *self-consistent* if the solution to the robust Ramsey policy is such that (i) $\Gamma_t = \Gamma_{-1}$ and (ii) the unconditional distribution of ρ_t for Φ_t is equal to ρ, for each $t \geq 0$. Our goal is to solve for a conditionally linear robustly optimal policy with a self-consistent initial commitment.[5]

Form the Lagrangian expression:

$$\begin{aligned} E_{-1} \sum_{t=0}^{\infty} \beta^t \Bigg\{ & - L(x_t, y_t, u_t) + \theta m_{t+1} \ln m_{t+1} - \phi_t (E_t m_{t+1} - 1) \\ & - \left[\mu'_{xt+1}, \mu'_{yt}\right] \left(\begin{bmatrix} x_{t+1} \\ E_t m_{t+1} y_{t+1} \end{bmatrix} - A \begin{bmatrix} x_t \\ y_t \end{bmatrix} - B u_t - C \varepsilon_{t+1} \right) \Bigg\}, \end{aligned}$$

where $\beta^t \phi_t$ and $\beta^t \left[\mu'_{xt+1}, \mu'_{yt}\right]$ are vectors of Lagrange multipliers associated with the constraints (1) and (8), respectively. Note that μ_{xt+1} is measurable with respect to date $t+1$ information and corresponds to x_{t+1}, but μ_{yt} is measurable with respect to date t information and corresponds to y_t.

The first-order conditions with respect to m_{t+1} are given by

$$\theta(1 + \ln m_{t+1}) - \phi_t - \mu'_{yt} y_{t+1} = 0. \tag{10}$$

Substituting the linear form in (9) into the above equation yields:

$$\ln m_{t+1} = -1 + \theta^{-1} \phi_t + \theta^{-1} \mu'_{yt} (\Phi_t + \Gamma_t \epsilon_{t+1}). \tag{11}$$

Since Γ_t, Φ_t, μ_{yt}, ϕ_t are measurable with respect to the date t information, m_{t+1} follows a *log-normal* distribution conditioned on date t information. Using equation (1), we can show

[5] See Woodford (2003, Chapter 7) for the concept of self-consistency under rational expectations.

that the worst-case density is given by

$$m_{t+1} = \exp\left(-\frac{1}{2}\theta^{-2}\mu_{yt}'\Gamma_t\Gamma_t'\mu_{yt} + \theta^{-1}\mu_{yt}'\Gamma_t\epsilon_{t+1}\right). \tag{12}$$

This implies that under the worst-case belief, the shock ε_t follows a normal distribution with mean $\theta^{-1}\Gamma_t'\mu_{yt}$ and covariance matrix I.

Under the worst-case belief, the conditional expectation of y_{t+1} is given by

$$E_t m_{t+1} y_{t+1} = \Phi_t + \theta^{-1}\Gamma_t\Gamma_t'\mu_{yt}, \tag{13}$$

and the conditional relative entropy is given by

$$E_t m_{t+1} \ln m_{t+1} = \frac{1}{2}\theta^{-2}\mu_{yt}'\Gamma_t\Gamma_t'\mu_{yt}. \tag{14}$$

Since $E_t y_{t+1} = \Phi_t$, it follows that, relative to the policymaker's expectations, the worst-case beliefs distort the private agents' expectations by $\theta^{-1}\Gamma_t\Gamma_t'\mu_{yt}$.

Now substitute equations (13) and (14) into the Lagrangian to obtain

$$\begin{aligned} E_{-1}\sum_{t=0}^{\infty}\beta^t \Bigg\{ & -L(x_t, \Phi_{t-1} + \Gamma_{t-1}\epsilon_t, u_t) + \frac{1}{2}\theta^{-1}\mu_{yt}'\Gamma_t\Gamma_t'\mu_{yt} \\ & - \left[\mu_{xt+1}', \mu_{yt}'\right]\left(\begin{bmatrix} x_{t+1} \\ \Phi_t + \theta^{-1}\Gamma_t\Gamma_t'\mu_{yt} \end{bmatrix} - A\begin{bmatrix} x_t \\ \Phi_{t-1} + \Gamma_{t-1}\epsilon_t \end{bmatrix} - Bu_t - C\varepsilon_{t+1}\right)\Bigg\} \end{aligned}$$

The first-order conditions with respect to x_t, u_t, Φ_t, and Γ_t are given by[6]

$$x_t: \quad 0 = -Q_{xx}x_t - Q_{xy}\left(\Phi_{t-1} + \Gamma_{t-1}\epsilon_t\right) - S_x u_t - \beta^{-1}\mu_{xt} + A_{xx}'E_t\mu_{xt+1} + A_{yx}'\mu_{yt}, \tag{15}$$

$$u_t: \quad 0 = -Ru_t - S_x'x_t - S_y'\left(\Phi_{t-1} + \Gamma_{t-1}\epsilon_t\right) + B_x'E_t\mu_{xt+1} + B_y'\mu_{yt}, \tag{16}$$

$$\Phi_t: \quad 0 = -\left(Q_{xy}'E_t x_{t+1} + Q_{yy}\Phi_t\right) - S_y E_t u_{t+1} - \beta^{-1}\mu_{yt} + E_t\left[A_{xy}'\mu_{xt+2} + A_{yy}'\mu_{yt+1}\right], \tag{17}$$

[6]For any scalar function $f(X)$ of a $n \times k$ matrix $X = (x_{ij})$, we define the derivative

$$\frac{\partial f(X)}{\partial X} = \begin{bmatrix} \frac{\partial f(X)}{\partial x_{11}} & \cdots & \frac{\partial f(X)}{\partial x_{1k}} \\ \cdots & & \\ \frac{\partial f(X)}{\partial x_{n1}} & \cdots & \frac{\partial f(X)}{\partial x_{nk}} \end{bmatrix}_{n\times k}.$$

$$\Gamma_t: \quad 0 = -\beta Q'_{xy} E_t x_{t+1} \epsilon'_{t+1} - \beta Q_{yy} E_t \left(\Phi_t + \Gamma_t \epsilon_{t+1}\right) \epsilon'_{t+1} - \beta S_y E_t u_{t+1} \epsilon'_{t+1}$$
$$- \theta^{-1} \mu_{yt} \mu'_{yt} \Gamma_t + \beta E_t \left[\left(A'_{xy} \mu_{xt+2} + A'_{yy} \mu_{yt+1}\right) \epsilon'_{t+1}\right]. \tag{18}$$

Plugging (9), (13) and (14) into (8) yields:

$$0 = x_{t+1} - A_{xx} x_t - A_{xy} \left(\Phi_{t-1} + \Gamma_{t-1} \epsilon_t\right) - B_x u_t - C_x \epsilon_{t+1}, \tag{19}$$

$$\theta^{-1} \Gamma_t \Gamma'_t \mu_{yt} + \Phi_t - A_{yx} x_t - A_{yy} \left(\Phi_{t-1} + \Gamma_{t-1} \epsilon_t\right) - B_y u_t = 0. \tag{20}$$

Since we focus on self-consistent optimal policy, we assume that $\Gamma_t = \Gamma$ for all t. We start with an initial guess of Γ and then solve for $\left\{x_t, u_t, \Phi_t, \mu_{xt}, \mu_{yt}\right\}$ given this guess. Equations (15), (16), (17), (19), and (20), together with two identity equations, $E_t \varepsilon_{t+1} = 0$ and $E_t \mu_{xt+1} = E_t \mu_{xt+1}$, form a linear system of the following form:

$$J \begin{bmatrix} E_t \varepsilon_{t+1} \\ E_t x_{t+1} \\ \Phi_t \\ E_t u_{t+1} \\ E_t \mu_{yt+1} \\ E_t \mu_{xt+1} \\ E_t \mu_{xt+2} \end{bmatrix} = F \begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \\ u_t \\ \mu_{yt} \\ \mu_{xt} \\ E_t \mu_{xt+1} \end{bmatrix}, \tag{21}$$

where

$$J = \begin{bmatrix} I & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & I & 0 \\ 0 & I & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & I & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & A'_{xx} & 0 \\ 0 & 0 & 0 & 0 & 0 & B'_x & 0 \\ 0 & Q'_{xy} & Q_{yy} & S_y & -A'_{yy} & 0 & -A'_{xy} \end{bmatrix},$$

and

$$F = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & I \\ A_{xy}\Gamma & A_{xx} & A_{xy} & B_x & 0 & 0 & 0 \\ A_{yy}\Gamma & A_{yx} & A_{yy} & B_y & -\theta^{-1}\Gamma\Gamma' & 0 & 0 \\ Q_{xy}\Gamma & Q_{xx} & Q_{xy} & S_x & -A'_{yx} & \beta^{-1}I & 0 \\ S'_y\Gamma & S'_x & S'_y & R & -B'_y & 0 & 0 \\ 0 & 0 & 0 & 0 & -\beta^{-1}I & 0 & 0 \end{bmatrix}.$$

Here ε_t, x_t and Φ_{t-1} are predetermined variables, and u_t, μ_{yt}, μ_{xt} and $E_t\mu_{xt+1}$ are non-predetermined variables.[7] Adapting a standard method for solving linear rational expectations models (e.g., Klein (2000)), we can show in Appendix A that the solution takes the following state space representation:[8]

$$\begin{bmatrix} \varepsilon_{t+1} \\ x_{t+1} \\ \Phi_t \end{bmatrix} = M \begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \end{bmatrix} + \begin{bmatrix} I \\ C_x \\ 0 \end{bmatrix} \varepsilon_{t+1}. \tag{22}$$

and

$$\begin{bmatrix} u_t \\ \mu_{yt} \\ \mu_{xt} \\ E_t\mu_{xt+1} \end{bmatrix} = N \begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \end{bmatrix}, \tag{23}$$

where the first row of M contains zero elements.

For Φ_t to have an invariant distribution, (22) should be a stationary process, which can be satisfied if M has all eigenvalues inside the unit circle. We assume this condition is satisfied.

After deriving the state space representation, we can derive an updated value for Γ using (18). When θ is finite, the presence of the term $\beta^{-1}\theta^{-1}\mu_{yt}\mu'_{yt}\Gamma_t$ in (18) implies that the solution for Γ may be time varying. The timeless perspective discussed in Woodford (2010) solves this issue. Since the optimal policy in the timeless perspective should satisfy (18) for every realization of the initial commitment, Φ_{-1}, we can take unconditional expectations on

[7] In our solution, the non-predetermined variable y_t in the equilibrium system from the private sector is rewritten using predetermined variables (Φ_{t-1}, ε_t). Thus, the Lagrange multiplier μ_{yt} associated with y_t becomes a non-predetermined variable.

[8] Note that the system in (21) does not fit exactly in Klein's (2000) general form.

both sides of (18) to derive [9]

$$\begin{aligned}\Gamma &= \beta\left(\theta^{-1}E\left[\mu_{yt}\mu_{yt}'\right]+\beta Q_{yy}\right)^{-1} \\ &\quad\times\left\{E\left[\left(A_{xy}'\mu_{xt+2}+A_{yy}'\mu_{yt+1}\right)\epsilon_{t+1}'\right]-Q_{xy}'C_x-S_yE\left[u_t\epsilon_t'\right]\right\},\end{aligned} \tag{24}$$

where we have assumed that the inverse exists. In Appendix B, we derive formulas for computing unconditional moments $E\left[\mu_{yt}\mu_{yt}'\right]$, $E\left[\mu_{xt+2}\epsilon_{t+1}'\right]$, $E\left[\mu_{yt+1}\epsilon_{t+1}'\right]$, and $E\left[u_t\epsilon_t'\right]$. The above equation gives an updated solution for Γ denoted by $\hat{\Gamma}$. If $\hat{\Gamma}$ is sufficiently close to Γ, stop. Otherwise, use $\hat{\Gamma}$ to replace Γ and repeat the above procedure until convergence. We are unable to give a condition to guarantee convergence. But this algorithm works well for all examples studied in this paper.

From the above analysis, we find that the robustness parameter θ appears explicitly in two terms only: one is $-\theta^{-1}\Gamma\Gamma'$ in the matrix F and the other is $\theta^{-1}E\left[\mu_{yt}\mu_{yt}'\right]$ in equation (24). In the limit as $\theta\to\infty$, both terms vanish and the solution converges to the rational expectations case. When θ is smaller, the impact of robustness is larger. In addition, the impact of robustness on Γ depends on the moment $E\left[\mu_{yt}\mu_{yt}'\right]$ for any fixed finite value of θ. These observations are important for understanding the intuition behind the examples analyzed later.

4. Robust Discretionary Policy

To understand the value and importance of policy commitment, we now study robust discretionary policy. In this case, the policymaker reoptimizes every period by taking the private agents' expectations about future policy as given. We shall focus on Markov perfect equilibrium (MPE). A robust MPE consists of time-invariant functions $f(x_t)$ and $V(x_t)$ such that given $y_{t+1}=f(x_{t+1})$, $V(x_t)$ and $y_t=f(x_t)$ solve the following Bellman equation:

$$V(x_t)=\max_{y_t,u_t}\min_{m_{t+1}}\quad -L(x_t,y_t,u_t)+\beta E_tV(x_{t+1})+\theta E_tm_{t+1}\ln m_{t+1}, \tag{25}$$

subject to (1) and (8).

We focus on linear solutions for f in that

$$y_t=Gx_t,\ \text{for all } t, \tag{26}$$

for some matrix G to be determined. Adapting the method from the rational expectations

[9] Note that taking unconditional expectations on other first-order conditions, i.e. (15)∼(17), does not add additional information to the solution.

analysis (e.g., Backus and Driffil (1986) and Soderlind (1999)), we use an iterative procedure to find G by backward induction. Specifically, let

$$y_{t+1} = G_{t+1} x_{t+1}, \tag{27}$$

where G_{t+1} is a deterministic matrix. Substituting this expression into the forward-looking equation in (8) and rewrite the Bellman equation (25) as

$$V_t(x_t) = \max_{y_t, u_t} \min_{m_{t+1}} \quad -L(x_t, y_t, u_t) + \beta E_t V_{t+1}(x_{t+1}) + \theta E_t m_{t+1} \ln m_{t+1}. \tag{28}$$

We can show that the first-order condition with respect to m_{t+1} is given by (10), where ϕ_t and μ_{yt} are the Lagrange multipliers associated with (1) and the forward-looking equation in (8), respectively. Using equations (1), (8), (10), (26) we can show that

$$m_{t+1} = \exp\left(-\frac{1}{2}\theta^{-2}\mu'_{yt} G_{t+1} C_x C'_x G'_{t+1} \mu_{yt} + \theta^{-1}\mu'_{yt} G_{t+1} C_x \epsilon_{t+1}\right).$$

It follows that m_{t+1} follows a *log-normal* distribution conditioned on date t information. We can then compute

$$E_t m_{t+1} \ln m_{t+1} = \frac{1}{2}\theta^{-1}\mu'_{yt}\Lambda_{t+1}\mu_{yt}, \tag{29}$$

$$E_t m_{t+1} y_{t+1} = G_{t+1}(A_{xx}x_t + A_{xy}y_t + B_x u_t) + \Lambda_{t+1}\mu_{yt}, \tag{30}$$

where we define $\Lambda_{t+1} \equiv \theta^{-1} G_{t+1} C_x C'_x G'_{t+1}$. Since $E_t y_{t+1} = G_{t+1}(A_{xx}x_t + A_{xy}y_t + B_x u_t)$, the worst-case beliefs distort the private agents' expectations by $\Lambda_{t+1}\mu_{yt}$, relative to the policy-maker's expectations.

Substituting (30) into (8) yields:

$$G_{t+1}(A_{xx}x_t + A_{xy}y_t + B_x u_t) + \Lambda_{t+1}\mu_{yt} = A_{yx}x_t + A_{yy}y_t + B_y u_t. \tag{31}$$

Conjecture that

$$V_t(x_t) = -\frac{1}{2}x'_t J_t x_t - \frac{1}{2}v_t,$$

where J_t is a deterministic matrix and v_t is a deterministic constant. Substituting the conjecture

and (29) into the Bellman equation (28) yields:

$$
\begin{aligned}
-\frac{1}{2}x_t' J_t x_t - \frac{1}{2}v_t \quad = \quad & \max_{y_t, u_t} \; -\frac{1}{2}x_t' Q_{xx} x_t - \frac{1}{2}y_t' Q_{yy} y_t - x_t' Q_{xy} y_t \\
& -\frac{1}{2}u_t' R u_t - x_t' S_x u_t - y_t' S_y u_t \\
& +\frac{1}{2}\mu_{yt}' \Lambda_{t+1} \mu_{yt} + \beta E_t \left(-\frac{1}{2}x_{t+1}' J_{t+1} x_{t+1} - \frac{1}{2}v_{t+1} \right),
\end{aligned} \tag{32}
$$

subject to (31) and

$$x_{t+1} = A_{xx} x_t + A_{xy} y_t + B_x u_t + C_x \varepsilon_{t+1},$$

where the Lagrange multiplier μ_{yt} is associated with (31). The above problem is a standard linear quadratic control problem with state variable x_t and control variables u_t and y_t. In Appendix C, we solve for J_t, v_t and linear decision rules of the following form:

$$u_t = -F_t x_t, \; y_t = G_t x_t.$$

To find the time-invariant functions, we start with an initial guess of a positive semi-definite matrix J_{t+1}, a constant v_{t+1}, and any matrix G_{t+1}. We then obtain the updated constant v_t and matrices J_t and G_t. Iterating this process until convergence gives J and G in a Markov perfect equilibrium. In addition, v_t converges to v in the value function and F_t converges to F, which gives the stationary policy rule. Note that since a MPE is a fixed point, there may be multiple equilibria or no equilibrium. For the simple New Keynesian model, Woodford (2010) provides explicit characterizations. But a similar characterization for our general model is not available.

5. Applications to Monetary Policy

In this section, we apply our general theory to the study of robustly optimal monetary policy in a New Keynesian model. The central bank's loss function is given by

$$E_{-1} \sum_{t=0}^{\infty} \beta^t \frac{1}{2} \left[\pi_t^2 + \lambda (x_t - x^*)^2 \right], \tag{33}$$

where π_t denotes the inflation rate, x_t denotes the output gap, and $x^* \geq 0$ denotes the target output gap. The value of x^* measures the degree of inefficiency of the steady state. If $x^* = 0$, the steady state is efficient. The parameter λ measures the weight to the output gap.

Empirical studies find evidence that inflation is persistent (e.g., Fuhrer (1996)). Thus, suppose that the New Keynesian Phillips curve (NKPC) is given by

$$\pi_t = \kappa x_t + \eta \pi_{t-1} + \beta(1-\eta) E_t m_{t+1} \pi_{t+1} + z_t, \tag{34}$$

where $\eta \in (0,1)$ measures the degree of backward-looking behavior exhibited by inflation.[10] Assume that $\kappa > 0$ and $\{z_t\}$ follows an AR(1) process:

$$z_t = \rho_z z_{t-1} + \sigma_z \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1), \tag{35}$$

where $\rho_z \in (0,1)$ and $\sigma_z > 0$. When $\eta = 0$, the model reduces to the basic New Keynesian model studied by Woodford (2010). When $\eta \in (0,1)$, Woodford's (2010) analytical method does not apply because the lagged inflation rate provides another state variable.

5.1. Timeless Perspective Commitment

The robustly optimal policy problem with commitment is given by

$$\max_{\{x_t,\pi_t\}} \min_{\{m_{t+1}\}} -E_{-1} \sum_{t=0}^{\infty} \beta^t \frac{1}{2} \left[\pi_t^2 + \lambda (x_t - x^*)^2\right] + \theta \mathrm{E}_{-1} \sum_{t=0}^{\infty} \beta^t m_{t+1} \ln m_{t+1} \tag{36}$$

subject to (34). To solve this problem using our general framework, we define a new predetermined state variable as $k_{t+1} \equiv \pi_t$ and rewrite (34) as

$$k_{t+1} = \pi_t, \tag{37}$$

$$E_t m_{t+1} \pi_{t+1} = \beta^{-1}(1-\eta)^{-1} \left(-\eta k_t + \pi_t - \kappa x_t - z_t\right). \tag{38}$$

To map this problem into our general framework, we note that k_t and z_t constitute the predetermined vector x_t, π_t is the forward-looking variable y_t, and x_t is the control variable u_t. Assume that the central bank commits to a policy in the form

$$\pi_{t+1} = \Phi_t + \Gamma \varepsilon_{t+1}. \tag{39}$$

Let $\mu_{\pi t}$ denote the Lagrange multiplier associated with (38). As in Section 3, we can show that the worst-case belief satisfies

$$E_t m_{t+1} \pi_{t+1} = \Phi_t + \theta^{-1} \Gamma^2 \mu_{\pi t}. \tag{40}$$

[10] See, e.g, Gali and Gertler (1999) for a microfounded model with this feature.

Since $E_t\pi_{t+1} = \Phi_t$, $\theta^{-1}\Gamma^2\mu_{\pi t}$ represents the distortion of the private agents' expectations relative to the central bank's. We can then use our algorithm developed in Section 3 to solve for a robustly optimal policy with timeless perspective commitment.

We set baseline parameter values for β, κ, λ, and x^* as in Woodford (2010): $\beta = 0.99$, $\kappa = 0.2$, $\lambda = 0.08$, $x^* = 0.05$.[11] We first set $\eta = 0$ in the top four panels of Figure 1. In particular, the first panel replicates the result in Woodford (2010) when $\rho_z = 0$. We confirm his result that when the central bank is more concerned for robustness, it is more conservative and inflation responds less to the cost-push shock in the sense that Γ decreases with $1/\theta$. In addition, Γ decreases more than proportionally with σ_z. This implies that *certainty-equivalence* does not apply to the robustly optimal case. Notably, Γ decreases with σ_z for large values of σ_z when ρ_z is sufficiently large. When the cost-push shock is sufficiently persistent as much as $\rho_z = 0.95$, robustly optimal Γ is almost independent of σ_z, as illustrated in the second panel of Figure 1. The middle two panels of this figure shows the impact of ρ_z for $\sigma_z = 0.02$ and 0.04. Under rational expectations, Γ is an increasing function of ρ_z, which can be easily verified analytically. When robustness matters, however, this relation does not apply. For low values of σ_z, e.g., $\sigma_z = 0.02$, Γ first increases with ρ_z and then decreases to zero. If the concern for robustness is very high, e.g., $\theta = 0.001$, Γ decreases with ρ_z for high σ_z, e.g., $\sigma_z = 0.04$. In the limiting case, $\Gamma = 0$ if the cost-push shock process is a unit root process (i.e., $\rho_z = 1$) for all $\theta < \infty$.

As Woodford (2010) points out, the central bank is not willing to let inflation increase in response to a positive cost-push shock when it is faced with a concern for robustness. The central bank fears that a large shock in inflation might affect the private agents' inflation expectations unfavorably to the central bank and thus the output-inflation trade-off might worsen if it allows inflation to increase. Naturally, the central bank's concern for the unfavorable change in the agent's forecast becomes larger when the cost-push shock is more persistent. As a result, the central bank is more conservative in setting the optimal inflation rate. If the shock is as highly persistent as a unit root process, the central bank commits to an inflation rate which does not depend on the current shock.

The bottom two panels of Figure 1 show the impact of inflation persistence η. For simplicity, we assume that the cost-push shock does not exhibit persistence, i.e., $\rho_z = 0$. In the rational expectations case (i.e., $\theta = \infty$), Γ increases with η until $\eta = 0.5$, then decreases. The shape of Γ shows a symmetry around $\eta = 0.5$. As the degree of concern for robustness increases,

[11] There is a typo in Woodford (2010). Parameter values of κ and x^* in his paper should be interchanged to replicate the figures in his paper.

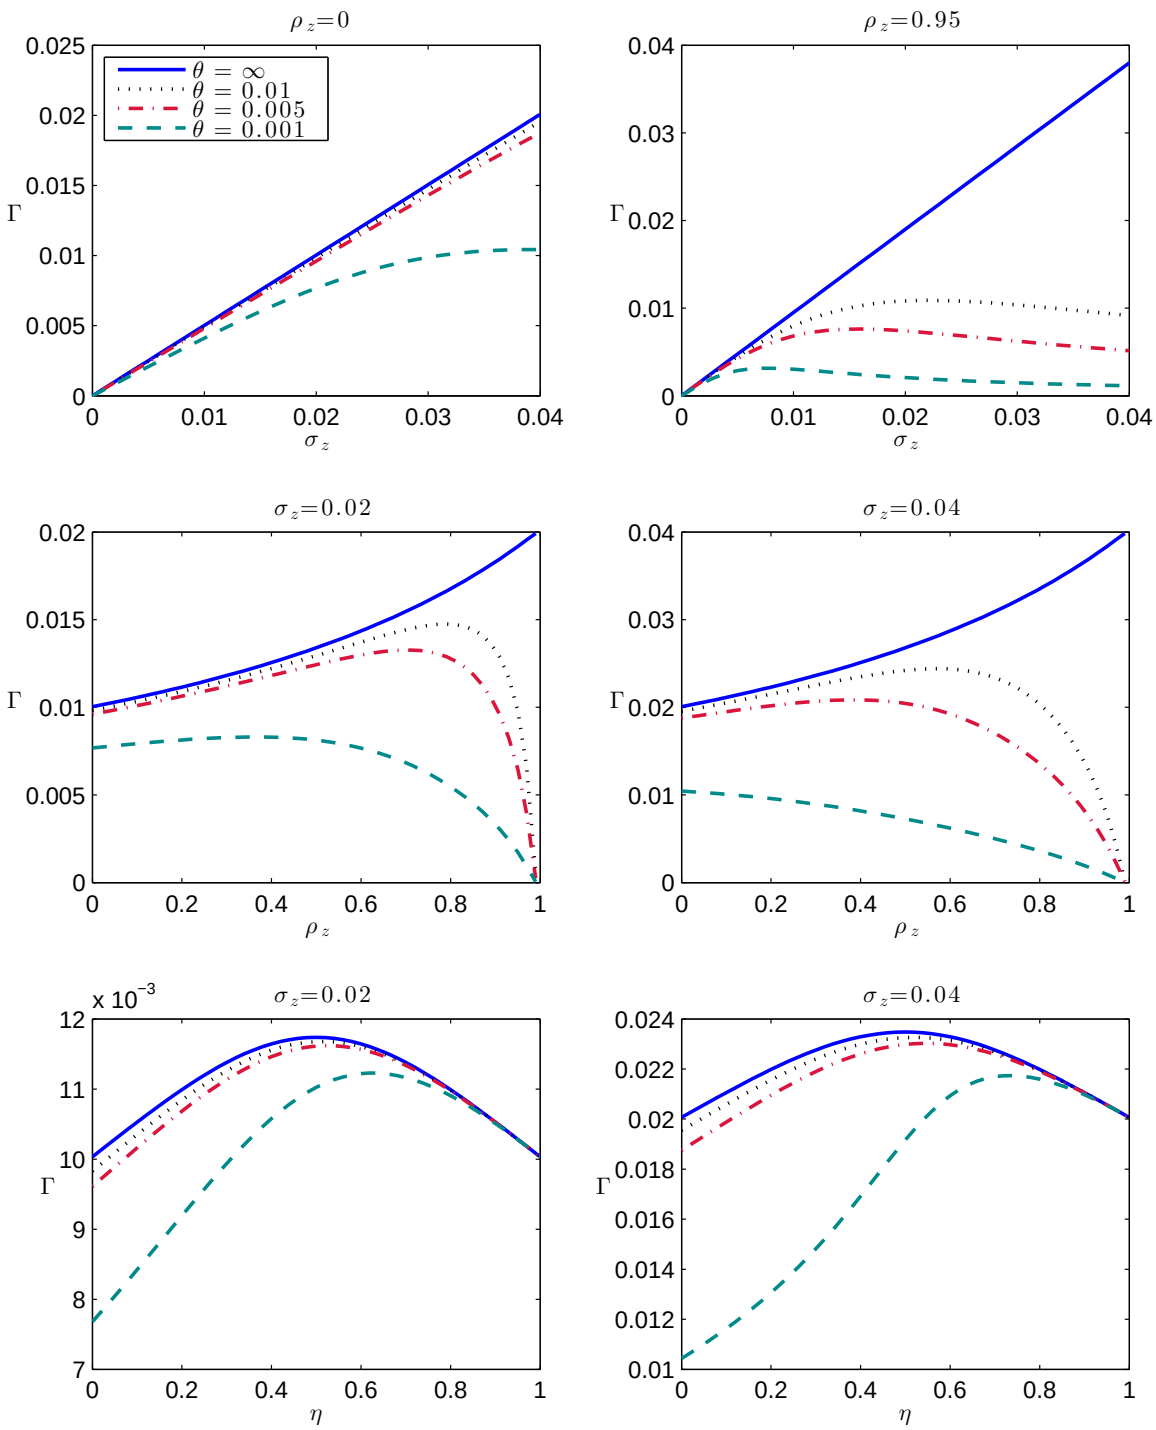

Figure 1: Variation of Γ with σ_z, ρ_z, and η for alternative values of θ. For the top four panels, we set $\eta = 0$. For the bottom two panels, we set $\rho_z = 0$.

i.e., θ decreases, Γ also decreases for a fixed η. The degree of decrease in Γ is larger when η is smaller. This is because the impact of the concern for robustness increases when the weight on the inflation expectation is larger (i.e., η is smaller) in the New Keynesian Phillips curve. In the limiting case when $\eta = 1$, inflation is completely determined one period ahead so that the central bank's doubt about the private sector's expectations does not change the optimal policy for inflation. Accordingly, the shape of Γ is no longer symmetric. Also note that for a given value of η the optimal value of Γ decreases more as σ_z increases.

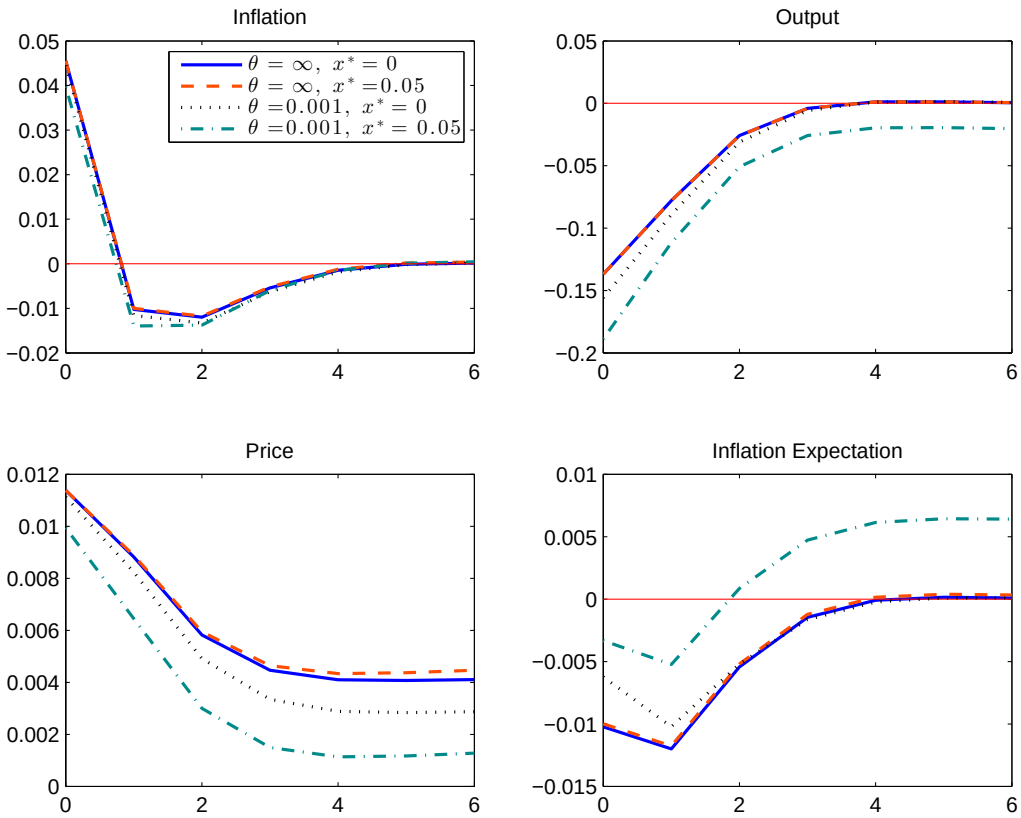

Figure 2: Impulse responses to a purely temporary positive cost-push shock with $\eta = 0.3$ and $\sigma_z = 0.02$. The inflation expectation is with respect to the worst-case beliefs of the private agents.

Figure 2 presents the dynamics of inflation, the worst-case inflation expectation, output gap, and prices in response to a purely temporary one-standard-deviation cost-push shock. For a moderate degree of inflation inertia ($\eta = 0.3$), the impulse responses exhibit very similar dynamics for both the rational expectations equilibrium and the robustly optimal equilibrium. The intuition is that the impact of the belief distortion from the private sector is small since a

weight of 0.7 is attached to the expected inflation. We also find that the long-run inflation rate is small but positive for both the rational expectations and robustness cases. The intuition is that in the presence of inflation inertia, the change in current inflation does not move future expected inflation one for one as it is attached a weight less than one. The central bank has an incentive to increase inflation to raise the output gap given $x^* > 0$. The robustly optimal long-run output gap is still negative, which follows from the fact the long-run worst-case inflation expectations of the private agents are positive, worsening the inflation-output tradeoff.

We emphasize that the impact of robustness also depends on the efficiency of the steady state, as reflected by the value of x^*. This is because the Lagrange multiplier $\mu_{\pi t}$ associated with NKPC (34) increases with x^* by the following first-order condition with respect to x_t,

$$\beta^{-1}(1-\eta)^{-1}\kappa\mu_{\pi t} = -\lambda(x_t - x^*).$$

As revealed by Figure 2, this result has two consequences. First, x^* biases the inflation expectation upward by (40). Thus, the steady-state output gap is negative even though the steady-state inflation rate is zero. Second, $E\left[\mu_{\pi t}^2\right]$ also increases with x^*. As we have discussed in the end of Section 3, this implies that Γ decreases with x^* for a finite θ (see (24)). Thus, under robustly optimal monetary policy the central bank will be more conservative in raising inflation in response to a positive cost-push shock when the economy is more inefficient. This result is in sharp contrast to that in the rational expectations case. Under rational expectations, the level of x^* does not affect Γ. It only affects the steady-state inflation.

5.2. Discretion

We define z_t and $k_t = \pi_{t-1}$ as the state variables and use the algorithm described in Section 4 to solve for a linear MPE of the form:

$$\pi_t = G_0 + G_\pi \pi_{t-1} + G_z z_t, \tag{41}$$

where G_0, G_π, and G_z are constants to be determined. The coefficients G_π and G_z measure the sensitivities of inflation to π_{t-1} and z_t , respectively. We can rewrite (41) as

$$\pi_t = G_0 + G_\pi \pi_{t-1} + \rho_z G_z z_{t-1} + \bar{G}\varepsilon_t,$$

where $\bar{G} \equiv \sigma_z G_z$, which is the sensitivity of inflation to a shock ε_t and corresponds to Γ under commitment.

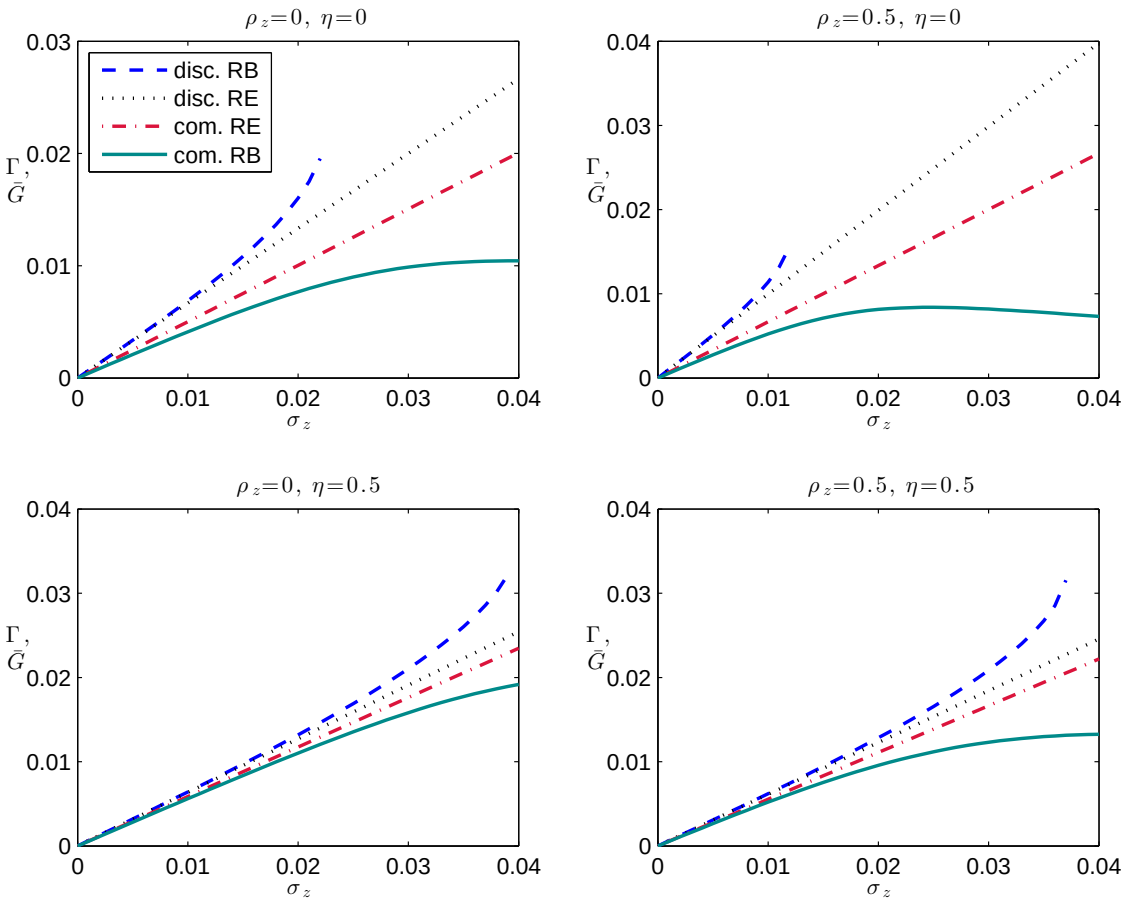

Figure 3: Variation of $\bar{G}$ and Γ with σ_z for different values of ρ_z and η, under discretionary policy and under an optimal commitment, with and without a concern for robustness ($\theta = \infty$ and $\theta = 0.001$).

As we showed earlier, the inflation inertia in the New Keynesian Phillips curve reduces the effect of inflation expectations and hence the effect of a concern for robustness on the optimal policy under commitment. This fact also applies to the discretionary policy with a concern for robustness. Figure 3 shows the optimal responses of inflation (i.e., Γ and $\bar{G}$) to a cost-push shock under commitment and under discretion for different values of σ_z, respectively, and the effects of η and ρ_z. For a purely temporary shock without inflation persistence (i.e., $\rho_z = 0$ and $\eta = 0$), we confirm Woodford's (2010) result that robustly optimal inflation under discretion is more responsive to the cost-push shock than in the rational expectations case. When the cost-push shock becomes more persistent, e.g., ρ_z is increased to 0.5, the robustly optimal inflation under discretion is more responsive to the cost-push shock for a given value of σ_z. But the range of σ_z for the existence of a robust MPE shrinks.

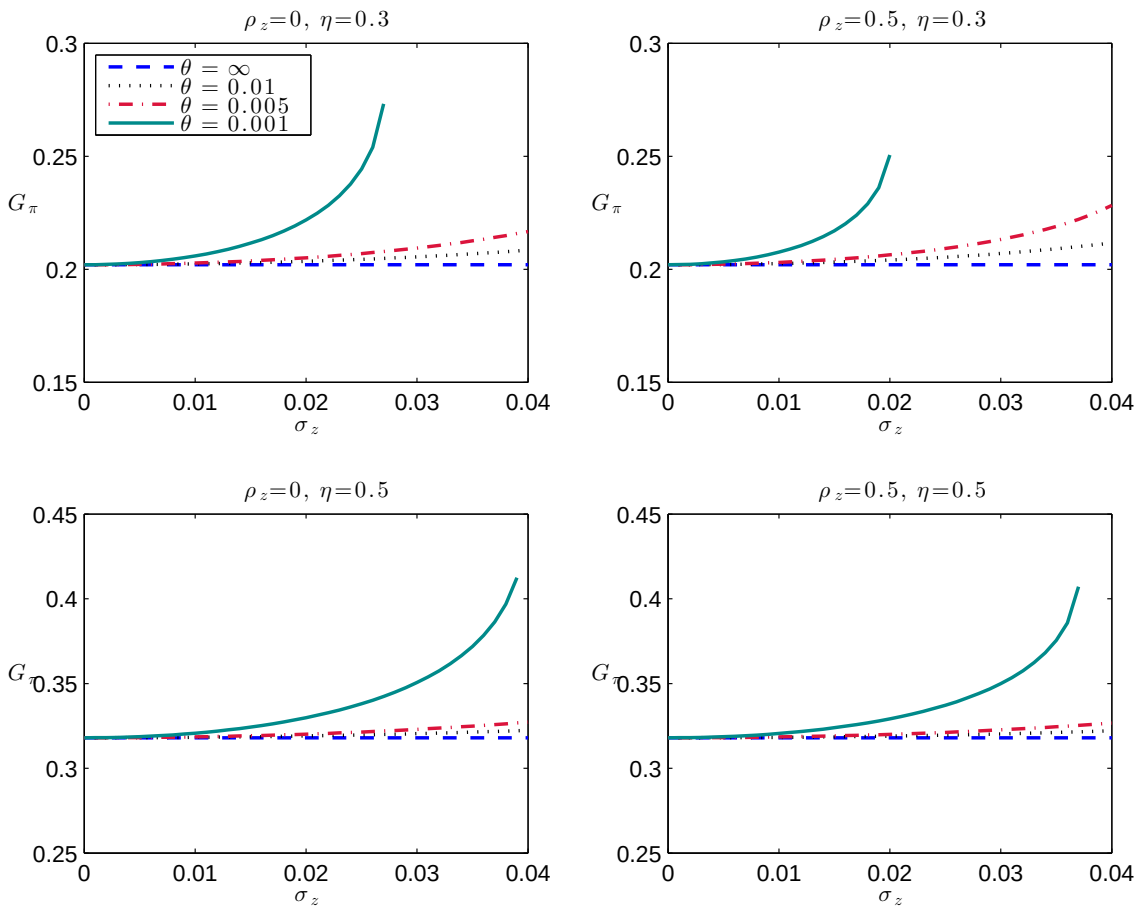

Figure 4: Variation of senstitvity of inflation to lagged inflation (G_π) with σ_z

Turn to the impact of inflation persistence. For $\rho_z = 0$, when the degree of inflation inertia is increased from $\eta = 0$ to $\eta = 0.5$, the robustly optimal inflation under discretion is less responsive to the cost-push shock for a given value of σ_z. In addition, the range of σ_z for the existence of a robust MPE expands. Thus, the net effects of the shock persistence and inflation persistence are ambiguous. Comparing the bottom two panels of Figure 3, we find that given $\eta = 0.5$, when ρ_z is increased from 0 to 0.5, the robustly optimal inflation under discretion becomes more responsive to the cost-push shock for small values of σ_z, but it is less responsive for high values of σ_z.

Figure 4 exhibits how a discretionary policy responds to lagged inflation for different values of σ_z. When the central bank has a doubt about its knowledge of the inflation expectations of the private sector, the optimal discretionary policy is more sensitive to lagged inflation than under rational expectations. The sensitivity increases as the volatility σ_z of the cost-push shock

increases. This is in contrast to the case of rational expectations under which optimal inflation responds to lagged inflation by the same degree regardless of the size of σ_z. Furthermore, as the central bank is more concerned about robustness (i.e., θ is smaller), the sensitivity to the lagged inflation of a discretionary policy becomes larger. But the range of σ_z for the existence of a robust MPE is smaller.

Figure 4 also reveals that for fixed ρ_z, when inflation peristence is larger (e.g., η is increased from 0.3 to 0.5), the robustly optimal inflation under discretion is more responsive to lagged inflation. For fixed η, when the cost-push shock is more persistent (e.g., ρ_z is increased from 0 to 0.5), the robustly optimal inflation under discretion is also more responsive to lagged inflation.

The intuition behind the above results is that, without commitment, the inflation bias is larger in the case of concerns for robustness than in the case of rational expectations. Under the worst-case beliefs, inflation expectations are biased upward, and hence a discretionary central bank has a greater incentive to respond to both the current shock and lagged inflation than under rational expectations. This effect is stronger if the cost-push shock is more persistent. However, this effect is weaker if a smaller weight is put on the inflation expectations.

Figure 5 shows the impulse responses of inflation, the worst-case inflation expectation, output gap, and prices to a positive cost-push shock under discretion. To compare with the commitment solution, we also plot impulse responses under commitment. When the central bank has a concern for robustness, the steady-state level of inflation under discretion is higher than that under rational expectations. But the steady-state level of the output gap under discretion is still below the one with commitment. With commitment, however, the steady-state level of inflation is not affected by a concern for robustness. This is because under robust MPE, the central bank needs to keep inflation rate higher since otherwise output would be lowered even further due to the distorted belief of private agents. But the central bank cannot allow too much inflation to increase output for fear of the possibility that inflation expectation would become so high that the output gap might actually decline further.

We emphasize that with discretionary monetary policy under robustness, the inflation expectation is distorted more than with commitment. With commitment, the central bank actively manages the beliefs of private agents and reduces the distortion in private beliefs; hence the output gap. By contrast, the central bank with discretionary monetary policy cannot manage expectations of the private sector effectively, and thus both the inflation expectation and the output gap worsen.

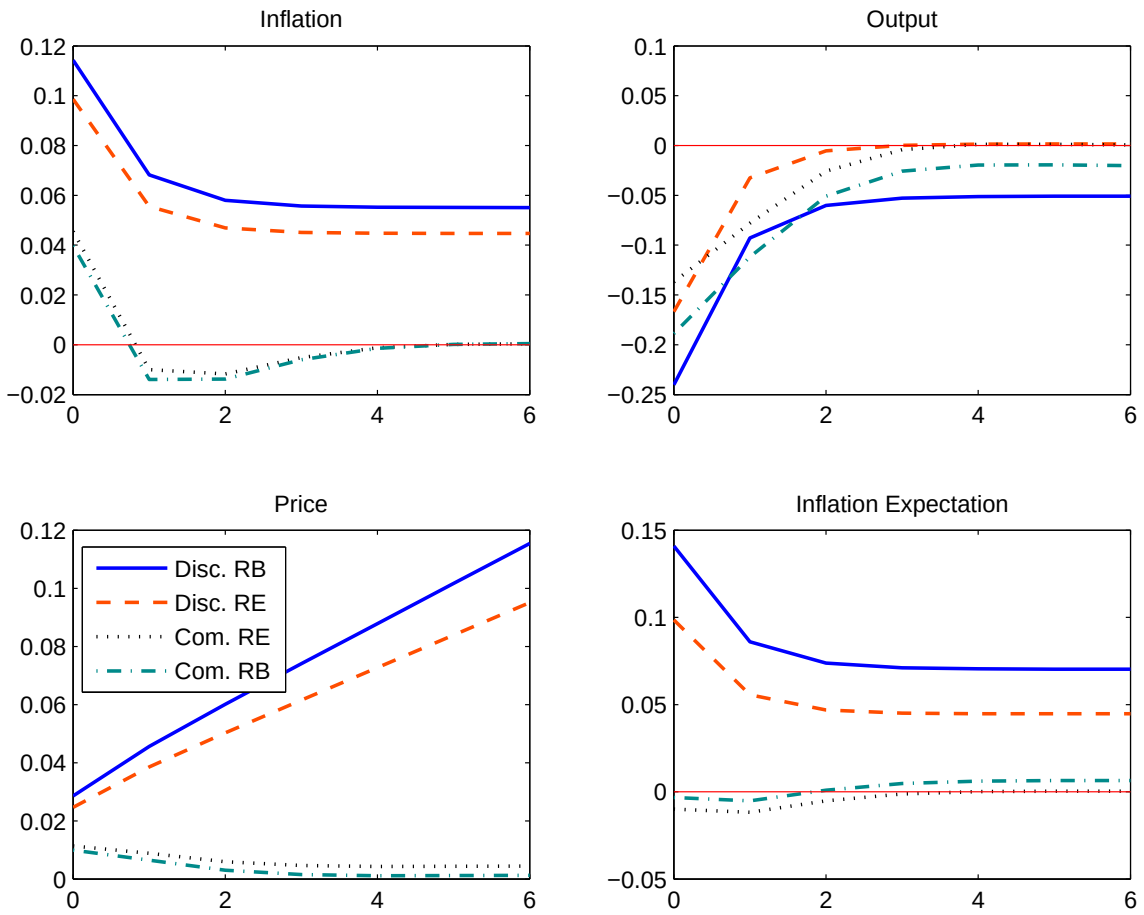

Figure 5: Impulse responses to a purely temporary positive cost-push shock with $\eta = 0.3$ and $\sigma_z = 0.02$ under commitment and discretion. With concern for robustness, $\theta = 0.001$.

6. Conclusion

In this paper, we have presented algorithms for solving robust policy problems in a general linear-quadratic framework by extending the Woodford (2010) approach. We have applied our algorithms to a New Keynesian model. Our algorithms are based on a linear-quadratic framework, which has been applied widely for policy analysis. One may wonder whether this framework can be justified by a microfounded model. This question does not seem obvious with concerns for robustness, though this issue is well understood under rational expectations (e.g., Woodford (2003)). Adam and Woodford (2012) present a fully microfounded model and show that the linear-quadratic model of Woodford (2010) can be microfounded.

This paper only focuses on the Woodford approach to robust policy. As Hansen and Sargent (2012) point out, there are other types of ambiguity for policy analysis. It would be interesting

to study these types of ambiguity in the general linear-quadratic framework. We leave this study to Kwon and Miao (2012).

Appendix

A State Space Representation of the Solution

We adapt Klein's (2000) method to derive the solution in (22) and (23). His method cannot be directly applied since the system in (21) does not fit in his general form. Apply a QZ decomposition to matrices J and F. There exist square complex matrices Q, S, T and Z such that

$$J = QSZ^H, \; F = QTZ^H,$$

where Z^H denotes the transpose of the complex conjugate of Z. Q and Z are unitary (i.e., $Q^HQ = Z^HZ = I$), and S and T are upper triangular (see Golub and van Loan (1996)). The decomposition can be reordered so that the block corresponding to the stable generalized eigenvalues (the ith diagonal element of T divided by the corresponding element in S) comes first. Define auxiliary variables

$$\begin{bmatrix} k_t \\ \lambda_t \end{bmatrix} = Z^H \begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \\ u_t \\ \mu_{yt} \\ \mu_{xt} \\ E_t\mu_{xt+1} \end{bmatrix}, \tag{A.1}$$

where k_t corresponds to the predetermined variables ε_t, x_t and Φ_{t-1}, and λ_t corresponds to the non-predetermined variables u_t, μ_{yt}, μ_{xt}, and $E_t\mu_{xt+1}$. Premultiplying (21) by Q^H, using (A.1), and partitioning S and T conformably, we can derive

$$\begin{bmatrix} S_{kk} & S_{k\lambda} \\ 0 & S_{\lambda\lambda} \end{bmatrix} E_t \begin{bmatrix} k_{t+1} \\ \lambda_{t+1} \end{bmatrix} = \begin{bmatrix} T_{kk} & T_{k\lambda} \\ 0 & T_{\lambda\lambda} \end{bmatrix} \begin{bmatrix} k_t \\ \lambda_t \end{bmatrix}. \tag{A.2}$$

We solve for a stable solution to the above system, which requires that the number of stable generalized eigenvalues be equal to the number of predetermined variables. We rule out the case of generalized eigenvalues with unitary modulus. Thus, the lower right block contains the unstable generalized eigenvalues. It follows that $\lambda_t = 0$ for all t. We can then write remaining

equations in the above system as

$$E_t k_{t+1} = S_{kk}^{-1} T_{kk} k_t, \tag{A.3}$$

where S_{kk} is invertible by our ordering of the matrix S.

Premultiplying both sides of (A.1) by Z yields:

$$\begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \\ u_t \\ \mu_{yt} \\ \mu_{xt} \\ E_t \mu_{xt+1} \end{bmatrix} = Z \begin{bmatrix} k_t \\ \lambda_t \end{bmatrix} = \begin{bmatrix} Z_{kk} & Z_{k\lambda} \\ Z_{\lambda k} & Z_{\lambda\lambda} \end{bmatrix} \begin{bmatrix} k_t \\ \lambda_t \end{bmatrix} = \begin{bmatrix} Z_{kk} \\ Z_{\lambda k} \end{bmatrix} k_t, \tag{A.4}$$

where we have used the fact that $\lambda_t = 0$. Suppose that Z_{kk} is invertible. It follows that

$$k_t = Z_{kk}^{-1} \begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \end{bmatrix}. \tag{A.5}$$

Substituting the above equation into (A.3) yields:

$$Z_{kk}^{-1} E_t \begin{bmatrix} \varepsilon_{t+1} \\ x_{t+1} \\ \Phi_t \end{bmatrix} = S_{kk}^{-1} T_{kk} Z_{kk}^{-1} \begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \end{bmatrix}.$$

Premultiplying by Z_{kk} and using the law of iterated expectations, we obtain

$$\begin{bmatrix} E_t \varepsilon_{t+1} \\ E_t x_{t+1} \\ \Phi_t \end{bmatrix} = Z_{kk} S_{kk}^{-1} T_{kk} Z_{kk}^{-1} \begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \end{bmatrix}.$$

Using equations (A.4) and (A.5) yields:

$$\begin{bmatrix} u_t \\ \mu_{yt} \\ \mu_{xt} \\ E_t \mu_{xt+1} \end{bmatrix} = Z_{\lambda k} k_t = Z_{\lambda k} Z_{kk}^{-1} \begin{bmatrix} \varepsilon_t \\ x_t \\ \Phi_{t-1} \end{bmatrix}.$$

Using (5), we then obtain the solution in the form of (22) and (23).

B Computing Unconditional Moments

After we derive the state space representation of the solution, we can compute unconditional moments of interest. First, define an auxiliary variable $X_t \equiv \begin{bmatrix} \varepsilon_t' & x_t' & \Phi_{t-1}' \end{bmatrix}'$. Then it follows from (22) that

$$E\left[X_{t+1}X_{t+1}'\right] = ME\left[X_tX_t'\right]M' + \begin{bmatrix} I & C_x' & 0 \\ C_x & C_xC_x' & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

This is a Lyapunov equation which determines the unconditional second moments of predetermined variables. This equation can be solved if predetermined variables are stationary processes, i.e. all the eigenvalues of M are inside the unit circle, which is assumed to be true under a timeless perspective.

Next, we use (23) to derive that

$$\mu_{yt} = N_{21}\varepsilon_t + N_{22}x_t + N_{23}\Phi_{t-1} = N_2X_t,$$

where $N_2 \equiv \begin{bmatrix} N_{21} & N_{22} & N_{23} \end{bmatrix}$ is the second row of N. We can compute that

$$E\left[\mu_{yt}\mu_{yt}'\right] = N_2E\left[X_tX_t'\right]N_2'.$$

Finally, we use (23) to compute the unconditional moments:

$$E\begin{bmatrix} u_{t+1}\varepsilon_{t+1}' \\ \mu_{yt+1}\varepsilon_{t+1}' \\ \mu_{xt+1}\varepsilon_{t+1}' \\ E_{t+1}\mu_{xt+2}\varepsilon_{t+1}' \end{bmatrix} = NE\begin{bmatrix} \varepsilon_{t+1}\varepsilon_{t+1}' \\ x_{t+1}\varepsilon_{t+1}' \\ \Phi_t\varepsilon_{t+1}' \end{bmatrix} = N\begin{bmatrix} I \\ C_x \\ 0 \end{bmatrix}.$$

We then obtain the moments $E\left[\mu_{yt}\mu_{yt}'\right]$, $E\left[\mu_{xt+2}\varepsilon_{t+1}'\right]$, $E\left[\mu_{yt+1}\varepsilon_{t+1}'\right]$, and $E\left[u_t\varepsilon_t'\right]$.

C Computing Discretionary Policy

Substituting $x_{t+1} = A_{xx}x_t + A_{xy}y_t + B_x u_t + C_x \varepsilon_{t+1}$ into the value function for x_{t+1} in (32) and deriving first-order conditions, we obtain

$$\begin{aligned} y_t : 0 = &- Q_{yy}y_t - Q'_{xy}x_t - S_y u_t + (A_{yy} - G_{t+1}A_{xy})' \mu_{yt} \\ &- \beta A'_{xy} J_{t+1} (A_{xx}x_t + A_{xy}y_t + B_x u_t), \end{aligned} \tag{C.1}$$

$$\begin{aligned} u_t : 0 = &- Ru_t - S'_x x_t - S'_y y_t + (B_y - G_{t+1}B_x)' \mu_{yt} \\ &- \beta B'_x J_{t+1} (A_{xx}x_t + A_{xy}y_t + B_x u_t). \end{aligned} \tag{C.2}$$

Define the matrix $\hat{B}_{t+1} \equiv (A_{yy} - G_{t+1}A_{xy})'$ and assume that it is invertible. Solving (C.1) for μ_{yt}, we can obtain

$$\mu_{yt} = \mathbf{A}_{t+1}x_t + \mathbf{B}_{t+1}y_t + \mathbf{C}_{t+1}u_t, \tag{C.3}$$

where

$$\begin{aligned} \mathbf{A}_{t+1} &\equiv \hat{B}_{t+1}^{-1} \left(Q'_{xy} + \beta A'_{xy} J_{t+1} A_{xx}\right), \\ \mathbf{B}_{t+1} &\equiv \hat{B}_{t+1}^{-1} \left(Q_{yy} + \beta A'_{xy} J_{t+1} A_{xy}\right), \\ \mathbf{C}_{t+1} &\equiv \hat{B}_{t+1}^{-1} \left(S_y + \beta A'_{xy} J_{t+1} B_x\right). \end{aligned}$$

Substituting (C.3) into (C.2) and solving for u_t will yield

$$u_t = \mathbf{A}^*_{t+1}x_t + \mathbf{B}^*_{t+1}y_t, \tag{C.4}$$

where

$$\begin{aligned} \mathbf{A}^*_{t+1} &\equiv -D_{t+1}^{-1} \left[S'_x - (B_y - G_{t+1}B_x)' \mathbf{A}_{t+1} + \beta B'_x J_{t+1} A_{xx}\right], \\ \mathbf{B}^*_{t+1} &\equiv -D_{t+1}^{-1} \left[S'_y - (B_y - G_{t+1}B_x)' \mathbf{B}_{t+1} + \beta B'_x J_{t+1} A_{xy}\right]. \end{aligned}$$

Note that we have defined

$$D_{t+1} \equiv R - (B_y - G_{t+1}B_x)' \mathbf{C}_{t+1} + \beta B'_x J_{t+1} B_x,$$

and assumed that it is invertible.

Substituting (C.4) into (C.3), we obtain

$$\mu_{yt} = \mathcal{A}_{t+1} x_t + \mathcal{B}_{t+1} y_t, \tag{C.5}$$

where

$$\mathcal{A}_{t+1} \equiv \mathbf{A}_{t+1} + \mathbf{C}_{t+1} \mathbf{A}^*_{t+1},$$
$$\mathcal{B}_{t+1} \equiv \mathbf{B}_{t+1} + \mathbf{C}_{t+1} \mathbf{B}^*_{t+1}$$

After plugging (C.4) and (C.5) into (31) and solving for y_t, we can derive the policy rule for period t as $y_t = G_t x_t$, where

$$G_t \equiv -\Delta_{t+1}^{-1} \left[G_{t+1} \left(A_{xx} + B_x \mathbf{A}^*_{t+1} \right) + \Lambda_{t+1} \mathcal{A}_{t+1} - A_{yx} - B_y \mathbf{A}^*_{t+1} \right],$$
$$\Delta_{t+1} \equiv G_{t+1} \left(A_{xy} + B_x \mathbf{B}^*_{t+1} \right) + \Lambda_{t+1} \mathcal{B}_{t+1} - A_{yy} - B_y \mathbf{B}^*_{t+1}.$$

We also have to assume that Δ_{t+1} is invertible. Substituting $y_t = G_t x_t$ into (C.4) yields:

$$u_t = -F_t x_t, \tag{C.6}$$

where

$$F_t \equiv -\mathbf{A}^*_{t+1} - \mathbf{B}^*_{t+1} G_t.$$

Applying the envelop theorem to (32) yields:

$$-J_t x_t = -Q_{xx} x_t - Q_{xy} y_t - S_x u_t + \hat{A}_{t+1} \mu_t - \beta A'_{xx} J_{t+1} \left(A_{xx} x_t + A_{xy} y_t + B_x u_t \right).$$

Substituting $y_t = G_t x_t$, (C.5), and (C.6) for y_t, μ_{yt} and u_t into the above equation, we can show that

$$\begin{aligned} J_t &= Q_{xx} + Q_{xy} G_t - S_x F_t - \hat{A}_{t+1} \left(\mathcal{A}_{t+1} + \mathcal{B}_{t+1} G_t \right) \\ &\quad + \beta A'_{xx} J_{t+1} \left(A_{xx} + A_{xy} G_t - B_x F_t \right). \end{aligned}$$

Matching constant terms in (32) yields:

$$v_t = \beta \left(\mathrm{tr} \left(C'_x J_{t+1} C_x \right) + v_{t+1} \right).$$

References

Adam, Klaus and Michael Woodford, 2011, Robustly Optimal Monetary Policy in a Micro-founded New Keynesian Model, Columbia University.

Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent, 2003, A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection, *Journal of the European Economic Association* 1, 68-123.

Backus, David, and John Driffil, 1986, The Consistency of Optimal Policy in Stochastic Rational Expectations Models, CEPR Discussion Paper 124.

Clarida, Richard, Jordi Gali, and Mark Gertler, 1999, The Science of Monetary Policy: A New Keynesian Perspective, *Journal of Economic Literature* 37, 1661-1707.

Dennis, Richard, 2008, Robust Control with Commitment: A Modification to Hansen-Sargent, *Journal of Economic Dynamics and Control* 32, 2061-2084.

Dennis, Richard, Kai Leitemo, and Ulf Soderstrom, 2009, Methods for Robust Control, *Journal of Economic Dynamics and Control* 33, 1604-1616.

Ellsberg, Daniel, 1961, Risk, Ambiguity and the Savage Axiom, *Quarterly Journal of Economics* 75, 643-669.

Epstein, Larry G. and Tan Wang, 1994, Intertemporal Asset Pricing under Knightian Uncertainty, *Econometrica* 62, 283-322.

Gali Jordi, and Mark Gertler, 1999, Inflation Dynamics: A Structural Econometric Analysis, *Journal of Monetary Economics* 44, 195–222.

Gilboa, Itzach and David Schmeidler, 1989, Maxmin Expected Utility with Non-unique Priors, *Journal of Mathematical Economics* 18, 141-153.

Giordani, Paolo and Paul Soderlind, 2004, Solution of Macromodels with Hansen-Sargent Robust Policies: Some Extensions, *Journal of Economic Dynamics and Control* 28, 2367-2397.

Golub, Gene H., Charles F. van Loan, 1996. *Matrix Computations.* The Johns Hopkins University Press, Baltimore and London.

Hansen, Lars Peter, 2007, Beliefs, Doubts and Learning: The Valuation of Macroeconomic Risk, *American Economic Review* 97, 1-30.

Hansen, Lars Peter and Thomas J. Sargent. 2001. Robust Control and Model Uncertainty, *American Economic Review* 91, 60-66.

—, 2008, *Robustness*, Princeton, New Jersey: Princeton University Press.

—, 2012, Three Types of Ambguity, forthcoming in *Journal of Monetary Economics*.

Ju, Nengjiu and Jianjun Miao, 2012, Ambiguity, Learning, and Asset Returns, *Econometrica* 80, 559-591.

Kwon, Hyosung and Jianjun Miao, 2012, Three Types of Robust Ramsey Problem in a Linear-Quadratic Framework, work in progress, Boston University.

Klein, Paul, 2000. Using the Generalized Schur Form to Solve a Multivariate Linear Rational Expectations Model, *Journal of Economic Dynamics and Control* 24, 1405-1423.

Leitemo, Kai and Ulf Soderstrom. 2008. Robust Monetary Policy in the New Keynesian Framework, *Macroeconomic Dynamics* 12 (Supplement S1), 126-135.

Sims, Christopher A, 2002, Solving Linear Rational Expectations Models, *Computational Economics* 20, 1-20.

Soderlind, Paul, 1999, Solution and Estimation of RE Macromodels with Optimal Policy, *European Economic Review* 43, 813–823.

Woodford, Michael, 2003, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton, NJ: Princeton University Press.

Woodford, Michael, 2010, Robustly Optimal Monetary Policy with Near-Rational Expectations, *American Economic Review* 100, 274-303.