Welfare-Maximizing Monetary Policy Under Parameter Uncertainty

Rochelle M. Edge, Thomas Laubach, and John C. Williams *

December 6, 2006

Abstract

This paper examines welfare-maximizing monetary policy in an estimated dynamic stochastic general equilibrium model of the U.S. economy where the policymaker faces uncertainty about the true values of model parameters. Uncertainty about parameters describing preferences and technology implies not only uncertainty about model dynamics but also uncertainty about the "natural" level of output that the central bank should aim to achieve. We analyze the characteristics and performance of alternative monetary policy rules given the estimated covariance of parameter estimates. We find that the natural rates of output is very imprecisely estimated. We then show that policy rules that rely on the output gap and therefore estimates estimates of the natural rate of output perform poorly under parameter uncertainty. Instead, optimal policies respond primarily to indirect signals regarding natural rates extracted from observed prices and wages.

JEL Codes:
Keywords: Long-run identifying restrictions, technology shocks, monetary policy, model uncertainty.

*Board of Governors of the Federal Reserve System, rochelle.m.edge@frb.gov; Board of Governors of the Federal Reserve System, tlaubach@frb.gov; and Federal Reserve Bank of San Francisco, john.c.williams@sf.frb.org (corresponding author). We thank seminar participants at the ECB, Humboldt Universität Berlin, and ..., for comments on this research project and paper. The views expressed herein are those of the authors and do not necessarily reflect those of the Board of Governors of the Federal Reserve System, its staff, or the management of the Federal Reserve Bank of San Francisco.

1 Introduction

This paper examines welfare-maximizing monetary policy in an estimated dynamic stochastic general equilibrium model of the U.S. economy where the central bank faces uncertainty about the true values of model parameters. In this framework, welfare is maximized when output and the real wage equal their "natural rates," i.e. their values absent nominal rigidities, at which point the real interest rate equals its natural rate. Owing to the presence of sticky prices and wages, this first-best outcome is not attainable and the central bank faces a tradeoff between minimizing deviations of output from its natural rate and minimizing fluctuations in price and wage inflation. The natural rates of output is therefore a key determinant of optimal monetary policy. But, in a micro-founded model, the natural rate of output is a function of the parameters describing preferences and technology, which are uncertain. More generally, as discussed by Levin and Williams (2005), parameter uncertainty implies not only uncertainty about model dynamics but also uncertainty about the objectives of policy.

We analyze the usefulness of measures of the natural rate of output in determining monetary policy when model parameters are uncertain. Giannoni (2002), Levin and Williams (2005), and Levin, Onatski, Williams and Williams (2005; henceforth LOWW) have explored aspects of monetary policy under parameter uncertainty in micro-founded models where the central bank aims to maximize household welfare, but these papers do not explicitly analyze the usefulness of natural rates as guides for policy. We use the estimated covariance of model parameters as a measure of parameter uncertainty. We find that the natural rate output is imprecisely estimated. We then show that policy rules that rely heavily on estimates of the output gap and therefore the natural rate of output as guides for monetary policy are problematic when model parameters and, by implication, natural rates are uncertain. Instead, under uncertainty, optimal policies rely more heavily on indirect signals regarding natural rates extracted from observed prices and wages than would be optimal if natural rates were known.

This paper contributes to the large literature that examines the usefulness of measures of the natural rates in the conduct of monetary policy (see Orphanides and Williams, 2002, and citations therein). This research has generally been conducted with an ad hoc policy objective and has treated movements in natural rates as exogenous and not directly related to parameter uncertainty. In such a setting, natural rate uncertainty is a form

of additive uncertainty and, under certain conditions, has no implications for the optimal monetary policy.[1] In contrast, we consider general parameter uncertainty that includes the "slopes" of macroeconomic relationships that affect the dynamic responses of natural rates to shocks. Thus, natural rate uncertainty is intrinsically connected to parameter uncertainty and certainty equivalence does not necessarily apply.

The remainder of the paper is organized as follows. Section 2 describes the model. Section 3 describes the estimation and reports the results. Section 4 examines optimal monetary policy assuming model parameters are known. Section 5 considers optimal policy under parameter uncertainty. Section 6 concludes.

2 The Model

In this section, we develop a fairly standard closed-economy model in which households choose consumption and set wages for their differentiated types of labor services, and in which firms produce using a CES aggregate of households' labor services as input and set prices for their differentiated products. We abstract from capital formation primarily because we would like to derive a tractable and intuitive approximation to households' preferences as the objective for monetary policy (which will be presented in the fourth section). The dynamics of nominal and real variables are determined by the resulting first-order conditions of optimizing agents. We allow for various frictions such as habit formation and adjustment costs that interfere with instantaneous full adjustment in response to shocks. We begin by presenting preferences and technology and then describe firms' and households' optimization problems. The log-linear system of equations used in the simulations is summarized in Appendix A.

[1] In particular, if the data generating process for the natural rates is known, then the separation principle and certainty equivalence applies. In that case, the optimal estimates of the natural rates are inserted into the optimal policy rule and the parameters of the optimal policy are unaffected by natural rate uncertainty. Certainty equivalence does not apply if the policy rule is not optimal but instead is described by a simple rule as in Orphanides and Williams (2002) and many other papers.

2.1 The production technology

The economy's final good, $Y_{f,t}$, is produced according to the Dixit-Stiglitz technology,

$$Y_{f,t} = \left(\int_0^1 Y_{f,t}(x)^{\frac{\theta_p - 1}{\theta_p}} dx \right)^{\frac{\theta_p}{\theta_p - 1}}, \tag{1}$$

where the variable $Y_{f,t}(x)$ denotes the quantity of the xth differentiated goods used in production and θ_p is the elasticity of substitution between the differentiated production inputs.

Final goods producers obtain their differentiated production inputs used in production from the economy's differentiated intermediate goods producers who supply an output $Y_{m,t}(x)$. Not all of the differentiated output produced by the intermediate goods producers is realized as inputs into final goods production; some is absorbed in price formulation. Specifically, the relationship between $Y_{f,t}(j)$ and $Y_{m,t}(j)$ is given by,

$$Y_{f,t}(j) = Y_{m,t}(j) - \frac{\chi_p}{2} \left(\frac{P_t(j)}{P_{t-1}(j)} - \Pi_{p,*} \right)^2 Y_{m,t}. \tag{2}$$

The second term in (2) denotes the cost of setting prices. This is quadratic in the difference between the actual change in price and the steady-state change in prices, $\Pi_{p,*}$.

The differentiated intermediate goods, $Y_{m,t}(j)$ for $j \in [0,1]$, are produced by combining each variety of the economy's differentiated labor inputs that are supplied to market activities (that is, $\{L_{y,t}(z)\}$ for $z \in [0,1]$). The composite bundle of labor, denoted $L_{y,t}$, that obtains from this aggregation implies, given the current level of technology A_t, the output of the differentiated goods, $Y_{m,t}$. Specifically, production is given by,

$$Y_{m,t}(j) = A_t L_{y,t}(j) \text{ where } L_{y,t}(j) = \left(\int_0^1 L_{y,t}(x,j)^{\frac{\theta_w - 1}{\theta_w}} dx \right)^{\frac{\theta_w}{\theta_w - 1}}. \tag{3}$$

where θ_w is the elasticity of substitution between the differentiated labor inputs. The log-level of technology, A_t, is modeled as a random walk with drift:

$$\ln A_t = \ln \Gamma_a + \ln A_{t-1} + \epsilon_t \tag{4}$$

where ϵ_t is an i.i.d. innovation, and Γ_a is the constant trend growth rate of technology. In this paper we restrict ourselves to permanent shocks to the level of technology, and normalize the steady-state growth rate Γ_a to 1.[2]

[2]In a companion paper, Edge, Laubach, and Williams (2004), we consider sector-specific shocks to technology and shifts in the growth rate of technology.

2.2 Preferences

Households derive utility from their purchases of the consumption good C_t and from their use of leisure time, equal to what remains of their time endowment $\bar{L}$ after $0 \leq L_{u,t}(i) \leq \bar{L}$ hours of labor are supplied to non-gratifying activities. Its preferences exhibit an endogenous additive habit (assumed to equal a fraction $\eta \in [0,1]$ of its consumption last period) and are nonseparable between consumption and leisure.[3] Specifically, preferences of household i are given by

$$E_0 \frac{1}{1-\sigma} \sum_{t=0}^{\infty} \beta^t \left[(C_t(i) - \eta C_{t-1}(i))(\bar{L} - L_{u,t}(i))^{\zeta} \right]^{1-\sigma}, \tag{5}$$

where β is the household's discount factor, and ζ is a measure of the utility of leisure. The economy's resource constraint implies that $\int_0^1 C_t(x)dx \leq Y_{f,t}$, where $Y_{f,t}$ denotes the output of the economy's final good.

Non-gratifying activities include supplying $L_{y,t}$ hours to the labor market and devoting time to setting wages. Consequenty, we define $L_{u,t}(i)$ as

$$L_{u,t}(i) = L_{y,t}(i) + \frac{\chi_w}{2} \left(\frac{W_t(i)}{W_{t-1}(i)} - \Pi_{w,*} \right)^2 L_{u,t}. \tag{6}$$

The second term in (6) denotes the cost of setting wages and is analogous to the cost of setting prices.

2.3 Firms' optimization problems

The final goods producing firm, taking as given the prices set by each intermediate-good producer for their differentiated output, $\{P_t(j)\}_{j=0}^1$, chooses intermediate inputs, $\{Y_{f,t}(j)\}_{j=0}^1$, so as to minimize the cost of producing its final output $Y_{f,t}$, subject its production technology, given by equation (1). Specifically, the competitive firm in each sector solves

$$\min_{\{Y_{f,t}(j)\}_{j=0}^1} \int_0^1 P_t(x) Y_{f,t}(x) dx \text{ s.t. } Y_{f,t} \leq \left(\int_0^1 Y_{f,t}(x)^{\frac{\theta_p - 1}{\theta_p}} dx \right)^{\frac{\theta_p}{\theta_p - 1}}. \tag{7}$$

This problem implies a demand function for each of the economy's intermediate goods given by $Y_{f,t}(j) = (P_t(j)/P_t)^{-\theta_p} Y_{f,t}$, where the variable P_t is the aggregate price level, defined by $P_t = (\int_0^1 (P_t(x))^{1-\theta_p} dx)^{\frac{1}{1-\theta_p}}$.

[3] Basu and Kimball (2002) argue that nonseparability between consumption and leisure has substantial empirical support

Each intermediate firm chooses the quantities of inputs (in this case labor) that it will use for production and the price that it will set for its output. It is convenient to consider these two decisions as separate problems. In the first step of the problem firm j, taking as given the wages $\{W_t(i)\}_{i=0}^1$ set by each household for its variety of labor, chooses $\{L_{y,t}(i,j)\}_{i=0}^1$ to minimize the cost of attaining the aggregate labor bundle $L_{y,t}(j)$ that it will ultimately need for production. Specifically, the materials firm j solves:

$$\min_{\{L_{y,t}(i,j)\}_{i=0}^1} \int_0^1 W_t(x)L_{y,t}(x,j)dx \text{ s.t. } Y_{m,t}(j) \leq A_t \left(\int_0^1 L_{y,t}(x,j)^{\frac{\theta_w - 1}{\theta_w}} dx \right)^{\frac{\theta_w}{\theta_w - 1}} \tag{8}$$

This cost-minimization problem implies that the economy-wide demand for type i labor is $L_{y,t}(i) = \int_0^1 L_{y,t}(i,x)dx = (W_t(i)/W_t)^{-\theta_w} (1/A_t) \int_0^1 Y_{m,t}(x)dx$ where W_t denotes the aggregate wage, defined by $W_t = (\int_0^1 (W_t(x))^{1-\theta_w} dx)^{\frac{1}{1-\theta_w}}$. The marginal cost function of producing the intermediate goods is $MC_t(j) = W_t/A_t$.

In setting its price, $P_t(j)$, the intermediate good producing firm takes into account the demand schedule for its output that it faces from the final goods sectors and the fact—as summarized in equation (2)—that by resetting its price it reduces the amount of its output that it can sell to final goods producers. The intermediate-good producing firm j, taking as given the marginal cost $MC_t(j)$ for producing $Y_{m,t}(j)$, the aggregate price level P_t, and aggregate final-goods demand $Y_{f,t}$, chooses its price $P_t(j)$ to maximize the present discounted value of its profits subject to the cost of re-setting its price and the demand curve it faces for its differentiated output. Specifically, the firm solves,

$$\max_{\{P_t(j)\}_{t=0}^\infty} E_0 \sum_{t=0}^\infty \beta^t \frac{\Lambda_{c,t}}{P_t} \{(1 + \varsigma_{\theta,p}) P_t(j) Y_{f,t}(j) - MC_t(j) Y_{m,t}(j)\}$$

subject to

$$Y_{f,t}(j) = Y_{m,t}(j) - \frac{\chi_p}{2} \left(\frac{P_t(j)}{P_{t-1}(j)} - \Pi_{p,*} \right)^2 Y_{m,t} \text{ and } Y_{f,t}(j) = \left(\frac{P_t(j)}{P_t} \right)^{-\theta_p} Y_{f,t}, \tag{9}$$

In (9) the discount factor that is relevant for discounting nominal revenues and costs between periods t and $t+j$ is $E_t \beta^j \frac{\Lambda_{c,t+j}/P_{t+j}}{\Lambda_{c,t}/P_t}$, where $\Lambda_{c,t}$ is the household's marginal utility of consumption in period t. The parameter $\varsigma_{\theta,p}$ is a subsidy that ensures that in the absence of nominal rigidities the model's equilibrium outcome is Pareto optimal.

Our choice of quadratic adjustment ocsts for modeling nominal rigidities contrasts with that of many other recent studies, which rely instead on staggered price-and wage-setting in the spirit of Calvo (1983) and Taylor (1980). We prefer the quadratic adjustment cost approach over staggered price- and wage-setting because the latter imply heterogeneity among

agents. Partly for this reason, models utilizing staggered price and wage setting typically assume that utility is separable between consumption and leisure, in which case perfect insurance among households against labor income risk eliminates heterogeneity of their spending decisions. By contrast, if wages are staggered and household utility is nonseparable, differences across households in labor supply (which will result due to differences in wages set) lead to differences across household in the marginal utility of consumption (and hence consumption), even if perfect insurance is able to equalize wealth across households. The quadratic adjustment cost model allows us to avoid heterogeneity across agents. In any case, the resulting price and wage inflation equations are very similar to those derived from Calvo-based setups with inertia as in Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2003).

2.4 Households' optimization problem

The household taking as given the expected path of the gross nominal interest rate R_t, the price level P_t, the aggregate wage rate W_t, its profits income, and its initial bond stock $B_{i,0}$, chooses its consumption $C_t(i)$ and its wage $W_t(i)$ to maximize its utility subject to its budget constraint, the cost of re-setting its wage, and the demand curve it faces for its differentiated labor. Specifically, the household solves:

$$\max_{\{C_t(i),W_t(i)\}_{t=0}^{\infty}} E_0 \frac{1}{1-\sigma} \sum_{t=0}^{\infty} \beta^t \left[(C_t(i) - \eta C_{t-1}(i))(\bar{L} - L_{u,t}(i))^{\varsigma} \right]^{1-\sigma}$$

subject to

$$E_t \left[\beta \frac{\Lambda_{c,t+1}/P_{c,t+1}}{\Lambda_{c,t}/P_{c,t}} B_{t+1}(i) \right] = B_t(i) + (1+\varsigma_{\theta,w}) W_t(i) L_{y,t}(i) + Profits_t(i) - P_t C_t(i),$$

$$L_{y,t}(i) = L_{u,t}(i) - \frac{\chi_w}{2} \left(\frac{W_t(i)}{W_{t-1}(i)} - \Pi_{w,*} \right)^2 L_{u,t}, \text{ and } L_{y,t}(i) = \left(\frac{W_t(i)}{W_t} \right)^{-\theta_w} \int_0^1 L_{y,t}(i,j) dj. \quad (10)$$

The parameter $\varsigma_{\theta,w}$ in the household's budget constraint is a subsidy that ensures that in the absence of nominal rigidities the model's equilibrium outcome is pareto optimal. The variable $B_t(i)$ in the budget constraint is the state-contingent value, in terms of the numeraire, of household i's asset holdings at the beginning of period t. We assume that there exists a risk-free one-period bond, which pays one unit of the numeraire in each state, and denote its yield—that is, the gross nominal interest rate between periods t and $t+1$—by $R_t \equiv \left(E_t \beta \frac{\Lambda_{c,t+1}/P_{t+1}}{\Lambda_{c,t}/P_t} \right)^{-1}$. Profits in the budget constraint are those rebated from firms, which are ultimately owned by households.

2.5 Natural rate variables

Our model has a counterpart in which all nominal rigidities are absent and prices and wages are fully flexible. In this model the cost minimization problems faced by the final goods producing firm and the intermediate goods producing firms continue to be given by equations (7) and (8). The intermediate goods producing firms' profit maximization problem is similar to equation (9) but with the price adjustment cost parameter χ_p set to zero. Likewise the households' utility maximization problem is given by equation (10) but with the wage adjustment cost parameter χ_w set to zero. We refer to the level of output and real one-period interest rate in this equilibrium as the natural rate of output, $\widetilde{Y}_t$, and interest, $\widetilde{R}_t$. We also define log deviations of these variables from their steady-state values, $\tilde{y}_t \equiv \log \widetilde{Y}_t - \log Y_*$ and $\tilde{r}_t \equiv \log \widetilde{R}_t - \log R_*$. These natural rates are functions of our model's structural shocks and are derived in Appendix A.

2.6 Monetary authority

In the model with nominal ridigities we assume that the central bank uses the short-term interest rate as its instrument in the spirit of the literature on monetary policy feedback reaction functions. This rate is determined in accordance with an interest rate feedback equation by which the short-term interest rate responds to deviations of price inflation from its steady-state level, $\pi_{p,t} \equiv \log \Pi_{p,t} - \log \Pi_{p,*}$, and to the output gap, $x_t \equiv \log Y_t - \log \widetilde{Y}_t$. We also allow for policy inertia by including the lagged short-term interest rate in the feedback equation. In particular, monetary policy is described by

$$r_t = \phi_r r_{t-1} + (1-\phi_r)\phi_\pi \pi_{p,t} + (1-\phi_r)\phi_x x_t + \epsilon_{r,t}, \tag{11}$$

where $r_t \equiv \log R_t - \log R_*$ and $\epsilon_{r,t}$ is an i.i.d. policy shock. Note that we have suppressed the constant that incorporates the steady-state levels of the interest and inflation rate.

2.7 Equilibrium

Our complete model consists of the first-order conditions (derived in Appendix A) describing firms' optimal choice of prices and households' optimal choices of consumption and wages, the production technology (3), the policy rule (11), the market clearing conditions $Y_t(j) = \int_0^1 C_{j,t}(i)di \ \forall j$ and $L_t(i) = \int_0^1 L_{i,t}(j)dj \ \forall i$, and the law of motion for aggregate technology (4). We now turn to the parametrization of our model.

3 Estimation

We estimate several of the structural parameters of our model using a minimum distance estimator. Specifically, we estimate a VAR on quarterly U.S. data using empirical counterparts to the theoretical variables in our model, and identify two of the model's structural shocks using identifying assumptions that are motivated by our theoretical model. We then choose model parameters so as to match the impulse responses to these two shocks implied by the model to those implied by the VAR.[4] Estimation of model parameters by impulse response function matching has sometimes been criticized for being *ad hoc* in selecting which properties of the data the model has to match. While we agree that it would also be interesting to explore full information estimation methods, we nevertheless think that the estimation undertaken here is valuable precisely because we focus on properties of the data that have a clear structural interpretation. A more serious issue is the potentially weak identification of several model parameters. Although Canova and Sala (2006) explore this problem specifically in the context of IRF matching, other estimation strategies are similarly susceptible to this problem. In this section we first describe the VAR and the identification of the two shocks, and then discuss our parameter estimates.

3.1 VAR specification and identification

The specification of our VAR is determined by the model developed in the previous section and our identification strategy for the structural shocks. Concerning the latter, we follow Galí (1999) and assume that the technology shock is the only shock that has a permanent effect on the level of output per hour. The monetary shock is identified by a standard restriction on contemporaneous responses. Our model and identifying assumptions combined suggest the inclusion of five variables in the VAR: the first difference of log output per hour, inflation, the log labor share, log hours per person, and the nominal funds rate. Output per hour, the labor share, and hours are the BLS measures for the nonfarm business sector, where the labor share is computed as output per hour times the deflator for nonfarm business output divided by compensation per hour.[5] Inflation is computed using the GDP

[4]Recent applications of this estimation strategy are Rotemberg and Woodford (1997), Amato and Laubach (2003), and Christiano, Eichenbaum, and Evans (2005).

[5]By contrast, Altig *et al.* (2002) and Galí, López-Salido, and Vallés (2003) compute labor productivity by dividing real GDP by total hours in the nonfarm business sector, which could be problematic because of

deflator. Population is the civilian population age 16 and over. Letting Y_t denote the vector of variables in the VAR, we view the data in the VAR as corresponding, up to constants, to the model variables

$$Y_t = [\Delta(y_t - l_t),\ \pi_t,\ y_t - l_t - w_t,\ \Delta l_t,\ r_t]' \tag{12}$$

where lower case letters denote logs of the model variables. We estimate the VAR over the sample 1966q2 to 2006q2, including four lags of each variable. Details of the implementation of the identification scheme are provided in Appendix B.

A potentially controversial aspect of our specification is the inclusion of hours per capita in first differences. Recent years have witnessed a vigorous debate among macroeconomists whether hours worked increase or decline following a technology shock. Francis and Ramey (2005) and Altig *et al.* (2002) have attributed differences in results among different studies to the issue whether hours per capita are included in levels – in which case the *level* of hours is usually found to rise immediately following a technology shock – or whether hours enter in first differences or some other detrended form – in which case the *level* of hours is often found to decline during the first few quarters following the shock. Our aim in this paper is not to decide this issue either way, but rather to treat the question of the initial response of hours to a technology shock as an important aspect of model uncertainty with respect to which a desirable monetary policy should be robust.

The dashed lines in the panels of Figure 1 show the impulse responses to a permanent one percent increase in the level of technology. The dashed-dotted lines present one-standard deviation bands around the impulse responses, computed by bootstrap methods.[6] Upon impact, output immediately rises about half-way to its new steady-state level, whereas hours worked decline by about 1/4 percent. Over the following eight quarters, output completes its adjustment while hours worked return to their original level. Interestingly, the responses of the real wage and inflation to a technology shock suggest a relative limited role for nominal rigidities, with the real wage rising closely in line with productivity, and inflation declining upon impact by almost a percentage point. The estimated response of

the trending share of real nonfarm business output in real GDP.

[6]To prevent the standard error bands from diverging over time, we discard draws for which the implied reduced-form VAR was estimated to be unstable, such as draws for which the largest eigenvalue of the coefficient matrix in the reduced form, written in companion form, exceeds .99. In total, about 14 percent of all draws are being rejected.

monetary policy is to accommodate the increase in output by keeping the real funds rate on balance unchanged.

Figure 2 shows the impulse responses of the variables to a one percentage point positive funds rate shock. The estimated responses of output, the real wage, and inflation to a funds rate shock are consistent with many studies on the effects of monetary policy. Output falls within three quarters by about 3/4 percent in response to one percentage point increase in the funds rate that takes eight quarters to die out. Hours decline closely in line with output, and the real wage falls. The response of inflation exhibits a price puzzle that lasts for two quarters; thereafter, inflation declines for six quarters, to about 0.3 percent below its original level. The responses to a funds rate shock are more precisely estimated than the responses to the technology shock.

3.2 Model parameter estimates

With the VAR impulse responses to a funds rate shock and a technology shock in hand, we proceed to estimate the structural and monetary policy parameters of our model. First, we calibrate four model parameters that have little effect on the dynamic responses to shocks. We set the discount factor, $\beta = 0.9924$, which corresponds to discounting the future at a 3 percent annual rate. We normalize the time endowment to unity. We set the steady-state rates of price and wage inflation to zero. The remaining parameters are estimated by minimizing the squared deviations of the responses of the five variables $[y_t,\ \pi_t,\ w_t,\ l_t,\ r_t]$ implied by our model from their VAR counterparts. The IRFs of these five variables in quarters 0 through 8 following a technology shock in quarter 0, and in quarters 1 through 8 following a funds rate shock (the response in the impact quarter being constrained by the identifying assumption) provide a total of 85 moments to match. The impulse responses of the five variables following the two identified shocks receive equal weights in our estimation.

Preliminary estimation results based on IRFs from the two identified shocks indicated that certain model parameters were at or very near theoretical lower bounds. The monetary policy response to the output gap was estimated to be slightly negative but near zero, perhaps because the theoretical notion of the output gap in our model bears little resemblance to measures of the output gap used by policymakers. We therefore impose that the coefficient ϕ_y in the policy rule (11) equals zero throughout. Estimates of the parameters measuring the degree of wage and price indexation were likewise at or near

Table 1: Parameter Estimates

Model Estimation						
	Both shocks		Technology shocks shock IRFs only		Monetary shocks shock IRFs only	
σ	2.54	(1.00)	1.20	(1.39)	4.05	(2.03)
ζ	5.41	(2.19)	9.91	(4.17)	3.05	(2.84)
η	0.60	(0.08)	0.64	(0.24)	0.50	(0.11)
κ_w	0.006	(0.000)	0.044	(0.042)	0.012	(0.004)
κ_p	0.124	(0.015)	0.140	(0.024)	0.008	(0.001)
ϕ_r	0.87	(0.00)	0.70	(0.026)	0.87	(0.00)
ϕ_π	1.00	(0.06)	1.00	(0.06)	1.00	(0.21)

Notes: Standard errors in parentheses. Estimates reported in the column labeled "baseline" are obtained by matching IRFs to both the technology and the funds rate shock. Estimates reported in the next two columns are obtained by matching only the IRFs to the technology or the funds rate shock.

zero, the theoretical lower bound, and we therefore imposed this value for both parameters in the estimation. In addition, because the parameters θ_w and χ_w appear only as a ratio in the linearized version of the model (see Appendix A), they are not separately identified; the same is the case for the parameters θ_p and χ_p. We therefore estimate the ratios $\kappa_w = (\theta_w - 1)(1 + \varsigma_{\theta,w})/(\chi_w \Pi_{w,*}) = \theta_w/(\chi_w \Pi_{w,*})$ and $\kappa_p = (\theta_p - 1)(1 + \varsigma_{\theta,p})/(\chi_p \Pi_{p,*}) = \theta_p/(\chi_p \Pi_{p,*})$. In the end, we estimated seven free parameters: $\{\sigma, \zeta, \eta, \kappa_w, \kappa_p, \phi_r, \phi_\pi\}$.

The estimated parameters for three different estimation strategies are shown in Table 1, with standard errors in parentheses. Standard errors are computed using the Jacobian matrix from the numerical optimization routine and the empirical estimate of the covariance matrix of the impulse responses from the bootstrap. The first column reports our "baseline" estimates, obtained by matching the impulse responses to both the technology and the funds rate shock. The next two columns report the results using IRF moments from only the technology shock or only the funds rate shock, respectively.

As discussed before, one feature of both sets of impulse responses is that real output and hours adjust gradually in response to the shocks. In the case of a permanent technology shock, Rotemberg and Woodford (1996) demonstrated that DSGE models without intrinsic

inertia will not display such hump-shaped patterns; instead, these variables jump on impact and adjust monotonically to their new steady-state values. We therefore find a significant role for habit persistence. Our estimate of the habit parameter η is close to the value estimated by Smets and Wouters (2003) and Christiano, Eichenbaum, and Evans (2005), but smaller than the estimate reported in Fuhrer (2000).

Conversely, the responses of real wages and inflation differ substantially depending on the source of the shock: rapid responses to technology shocks, and sluggish to funds rate shocks. As we will discuss further below, this is a feature that our price and wage specification cannot deliver. While the small estimate of κ_w suggests that nominal wages are sticky, κ_p is large whenever we match the IRFs to the technology shock. As is clear from Figures 1 and 2, the model parameters are geared towards matching the responses to the technology shock, not the funds rate shock.[7]

Our estimates of the parameters of the monetary policy rule, ϕ_r and ϕ_π, are broadly consistent with the findings of many other studies that estimate monetary policy reaction functions, such as that of Clarida, Galí, and Gertler (2000).

Given the differential responses of several variables to the two different shocks, it is not surprising that the parameter estimates are sensitive to which set of IRFs is being matched. We explore this dimension of uncertainty about the true parameter values by reporting in the second two columns the parameter estimates that help the model match as closely as possible the IRFs to a technology shock, and in the final two columns the estimates that obtain when matching the funds rate shock IRFs only. The estimate of σ shown in the middle column is smaller than its baseline values, enabling the response of output to a technology shock to converge more rapidly to its steady-state value. This comes at the expense of not matching the protracted decline of output in response to a funds rate shock. The estimated degree of habit formation, and the parameters describing wage and price setting, are similar to the estimates matching both sets of IRFs.

When matching only the IRFs to the funds rate shock, we find a considerably higher estimate of σ compared to the baseline, whereas the estimates of ζ and η are slightly lower. The higher value of σ enables the model to match the gradual output response to the

[7]This result may depend in part on the fact that we are weighting all differences between model and VAR responses equally, instead of inversely proportional to the standard error around the VAR responses. Because the wage and inflation responses to a technology shock are larger in absolute terms than those to a funds rate shock, the estimates are geared towards matching the former more closely.

interest rate increase. The most striking change occurs for the parameter κ_p characterizing the degree of price stickiness. The much higher degree of price rigidity is necessary to match the gradual response of inflation to the funds rate shock.

In summary, even though the standard errors suggest that some parameters are estimated precisely, at least four of the five structural parameters (σ, ζ, and the two wage and price parameters) are either estimated imprecisely, or depend very sensitively on the choice of moments to be matched. In the following sections we explore the consequences of this source of uncertainty for optimal monetary policy.

4 Optimal Monetary Policy without Parameter Uncertainty

In this section we compute the optimal policy response to a technology shock assuming all parameters are known. We assume that the central bank objective is to maximize the unconditional expectation of the welfare of the representative household. We further assume that the central bank has the ability to commit to future policy actions; that is, we examine optimal policy under commitment, as opposed to discretion.

4.1 Approximating Household Welfare

We approximate household utility with a second-order Taylor expansion around the deterministic steady state. We denote steady-state values with an asterisk subscript. As shown in the Appendix, the second-order approximation of the period utility function depends on the squared output gaps (the log difference between output and its natural rate), the quasi-difference of the output gap, the cross-product of the output gap and its quasi-difference, and the squared price and wage inflation rates. Denote the log of the natural rate of output, that is, the log of the level of output that would obtain absent nominal frictions, by $\tilde{y}_t$. As shown in the appendix, in the linearized model, the natural rate of output is a function of leads and lags of the technology shock. Denote the output gap by $x_t = y_t - \tilde{y}_t$.

After numerous substitutions, the second-order approximation to utility can be written as follows:

$$\frac{1}{1-\sigma} \cdot \frac{\left(C_t - \eta C_{t-1}\right)^{1-\sigma} \left(\bar{L} - L_{u,t}\right)^{\zeta(1-\sigma)}}{\left(C_* - \eta C_*\right)^{1-\sigma} \left(\bar{L} - L_{u,*}\right)^{\zeta(1-\sigma)}}$$

$$
\begin{aligned}
= \quad & -\frac{1}{2}\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\left(\frac{1-\beta\eta}{1-\eta}\right)^2 x_t^2 \\
& -\frac{1}{2}\cdot\frac{\sigma}{(1-\eta)^2}\cdot\Big(x_t-\eta x_{t-1}\Big)^2 \\
& -(1-\sigma)\cdot\frac{1-\beta\eta}{(1-\eta)^2}x_t\Big(x_t-\eta x_{t-1}\Big) \\
& -\frac{1}{2}\cdot\frac{1-\beta\eta}{1-\eta}\left\{\ \chi_p\cdot\Pi_{p,*}^2\pi_{p,t}^2\ +\chi_w\cdot\Pi_{w,*}^2\pi_{w,t}^2\right\} \\
& +\text{T.I.P.}
\end{aligned}
$$

The first three right-hand side terms correspond to the welfare costs associated with output deviating from its natural rate. Owing to the presence of habit formation, both the level of the output gap and its quasi-difference affect welfare. Note that all three preference parameters enter in the coefficients of the welfare loss for these three terms. The final term corresponds to the welfare costs associated with adjustment costs in changing prices and wages. The coefficients in these terms depend primarily on the parameters associated with nominal rigidities. Importantly, the welfare costs of sticky prices and wages depend on the inverses of the price and wage sensitivity parameters, κ_p and κ_w, respectively. The more flexible are prices, the smaller are the welfare costs associated with a given magnitude of the inflation gap, and similarly for wages.

In the following, it is useful to report the welfare loss given by the left-hand side of the equation above, which we denote by $\mathcal{L}$. We also report the three components of the normalized loss, $\mathcal{L}_y, \mathcal{L}_p$, and $\mathcal{L}_w$, that correspond to the components of the welfare loss associated with output gaps, price inflation, and wage inflation, respectively.

4.2 Optimal Monetary Policy

To compute the optimal policy for a given set of parameter values, ignoring parameter uncertainty, we maximize the quadratic approximation of welfare subject to the constraints implied by the linearized model. We compute the fully optimal policy using Lagrangian methods as described in Finan and Tetlow (2001). We assume that the technology shock is the only stochastic element in the model and that it has unit variance. For these computations, we assume a discount rate arbitrarily close to zero, so that we are maximizing the unconditional measure of welfare. We repeat this exercise for each estimated version of the model; the results are shown in Table 2. The upper part of the table shows the welfare loss under the optimal policy, and the breakdown into its component parts (the components

Table 2: Optimal Monetary Policy without Parameter Uncertainty

	Model Estimation		
	Both shocks	Technology shocks	Monetary shocks
Welfare Losses			
$\mathcal{L}$	7.530	6.175	18.918
$\mathcal{L}_x$	.005	.037	.343
$\mathcal{L}_p$	7.232	5.120	8.496
$\mathcal{L}_w$	.292	1.018	10.079
Standard deviations			
x_t	.028	.133	.150
π_p	.402	.358	.108
π_w	.017	.089	.146

add to the total welfare loss, subject to rounding). The lower part of the table reports the resulting unconditional standard deviations of the output gap and price and wage inflation rates.

The total welfare costs under the optimal policy differ for the three parameterizations of the model, primarily reflecting differences in the stickiness in prices and wages across the estimated models. The welfare costs are the highest in the model estimated using monetary shocks, reflecting the very high degree of price and wage stickiness in that model. Similarly, the welfare costs are lowest in the model estimated using only technology shocks, in which prices and wages are relatively flexible.

In conducting our analysis of policy under parameter uncertainty in the next section it is useful to describe policy in terms of a so-called "instrument rule" where the short-term interest rate is determined by a small number of observable variables. As has been noted in the literature, the optimal policy rule can be implemented in a variety of ways, some of which may have different properties under uncertainty than others. We exclude from the determinants of policy expectations of future variables because these would implicitly require knowledge of model parameters. We consider a simple Taylor-type monetary policy rule where the nominal interest rate is determined by the price inflation rate, the output

Table 3: Optimized Simple Monetary Policy Rules without Parameter Uncertainty

	Model Estimation		
	Both shocks	Technology shocks	Monetary shocks
Welfare Losses			
$\mathcal{L}$	7.530	6.325	19.231
$\mathcal{L}_y$	.005	.000	.142
$\mathcal{L}_p$	7.233	4.817	7.738
$\mathcal{L}_w$	.291	1.509	11.351
Standard deviations			
x_t	.028	.007	.093
π_p	.402	.368	.103
π_w	.017	.035	.155
Coefficients			
ϕ_y	1.55	11.95	4.14
ϕ_p	1.40	1.00	1.00
ϕ_w	–	–	–

gap, and, in some cases, the wage inflation rate:

$$r_t = \bar{r} + \phi_p \, \pi_{p,t} + \phi_x \, x_t + \phi_w \, \pi_{w,t}, \tag{13}$$

where $\bar{r}$ is the constant long-run real interest rate, which is assumed to be known by the policymaker.

We start by considering a standard Taylor-type rule with no response to wages, i.e., $\phi_w = 0$. We optimized the coefficients of this rule to maximize unconditional welfare of the representative household using a hill-climber routine, using numerical methods described in Levin, Wieland, and Williams (1999). We imposed the restriction that the coefficient on price inflation exceeds unity. The results from this exercise are reported in Table 3.

In the benchmark model estimated using both shocks, the optimized simple rule is able to replicate the outcomes under the fully optimal policy. This rule features moderately strong responses to both the price inflation rate and the output gap. In the other two estimated models, the optimized simple rules are characterized by a minimal response to inflation and relatively large responses to the output gap. In both of these cases where the model is estimated based on one shock, these rules come very close, but do not quite

match, the welfare under the fully optimal rules.[8] Because this class of rules is able to nearly mimic the fully optimal policies in all three estimated versions of the model, we use it as our benchmark specification in the analysis of optimal policy under uncertainty in the following section.

5 Monetary Policy under Parameter Uncertainty

In this section, we analyze the performance and robustness of various monetary policies under parameter uncertainty. We consider two aspects of parameter uncertainty. One is due to sample variation and is measured by the covariance of the estimated parameters. The second is due to uncertainty regarding the identification assumptions used in the estimation of the model. We have two identified shocks that are used in estimating the model. If both identification assumptions are valid, then the estimates using both shocks will one average be more accurate. But, if the identification assumptions regarding one of the shocks is invalid, then the parameter estimates using those shocks are inconsistent. Given this uncertainty about estimation strategies, the uncertainty regarding the parameters may be greater than that implied by any specific set of estimates.

Throughout the following, we assume that the policymaker is uncertain about the five structural parameters and does not learn over time. For a given policy rule, the expected welfare loss is computed integrating over the distribution for these five parameters. Note that in these calculations, we fully take into account the effects of parameter values on the parameters of the loss function as in Levin and Williams (2005).

5.1 Uncertainty about the Natural Rate of Output

Before proceeding with the analysis of policy rules, we first provide some measure of the degree of uncertainty regarding the natural rate of output owing to parameter uncertainty. The response of the natural rate of output to a technology shock depends on three parameters describing household preferences: σ, η, and ζ. We use the estimated covariance matrix of the parameter values as a measure of this uncertainty. We assume the parameters are

[8] We repeated the analysis with rules that also respond to wage inflation. These three-parameter rules did modestly better and came closer to the fully optimal policies in the models estimated from IRFs to only a single source of shocks. In both of those cases, the optimal responses to wage inflation and the output gap are very large, around 20-30.

jointly normal distributed. We approximate the distribution with 1000 sets of parameters drawn randomly from the covariance matrix. Where necessary, we truncate the parameter values draws so that they lie within the allowable range of values. This approach is close to that followed in LOWW (2005), except that they apply Bayesian estimation techniques and therefore the posterior distribution of the parameter estimates is not normal.

The response of the natural rate of output to a technology shock can be highly uncertain. The thick solid line in the upper left panel of Figure 3 plots the impulse response of the natural rate of output to a one percentage point positive technology shock based on the point estimates from the model estimated using the IRFs from the technology shock alone. The thin solid line shows the median response from the 1000 draws from the parameter distribution; the dashed-dotted lines show the boundaries of the 70 and 90 percent confidence bands of the model impulse responses. Note that the structure of the model implies that there is no uncertainty about the long-run effects of technology shocks on the natural rate of output, which increases one-for-one with the level of technology. In this estimated version of the model, the preference parameters are estimated with little precision and, as a result, the model response of the natural rate of output varies widely. The uncertainty about the natural rate of output is less severe in the other two estimated versions of the model, but still substantial. The second and third panels show the distributions of the impulse responses of the natural rate of output for the other sets of model estimates, as indicated in the labels of the figure. The bottom right panel combines the impulse responses of the three sets of estimates and reports the resulting "combined" distribution that accounts for parameter uncertainty within a given estimation strategy and the uncertainty regarding the estimation strategies. In constructing this combined distribution, the three model estimation strategies are given equal weight.

Uncertainty about the natural rate of output implies a corresponding uncertainty about the natural rate of interest. Figure 4 repeats the plots of Figure 3 but this time showing the responses of the natural rate of interest. This figure shows that parameter uncertainty implies a great deal of uncertainty about the magnitude, and even the direction real interest rates need to respond to keep output at its natural level.

Table 4: Performance of Benchmark Policies under Parameter Uncertainty

	Model Estimation								
	Both shocks			Tech shocks			Monetary shocks		
	Opt.	Simple Rule		Opt.	Simple Rule		Opt.	Simple Rule	
Known $\tilde{y}$?	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No
Welfare Losses									
$\mathcal{L}$	7.554	7.646	8.015	10.132	52.195	5442	19.089	20.074	22.573
$\mathcal{L}_y$	.006	.082	.378	.032	.096	11	.306	.632	2.674
$\mathcal{L}_p$	7.258	7.242	7.261	9.136	8.578	294	9.563	8.640	9.558
$\mathcal{L}_w$	.292	.323	.376	.964	43.521	5137	9.220	10.802	10.341
Standard deviations									
x	.023	.037	.080	.066	.011	.161	.110	.096	.168
π_p	.402	.401	.402	.360	.350	.697	.111	.106	.111
π_w	.017	.018	.018	.086	.111	.493	.139	.150	.147
Policy rule coefficients									
ϕ_p	–	1.55	1.55	–	1.00	1.00	–	1.00	1.00
ϕ_y	–	1.40	1.40	–	11.95	11.95	–	4.14	4.14

5.2 Performance of Benchmark Rules under Parameter Uncertainty

In order to provide a benchmark for policies under uncertainty, we first compute the optimal outcome if the policymaker knew all the parameter values and followed the fully optimal policy in each case. We average the outcomes and losses over the 1000 draws of the parameters and report the results in the first column of Table 4. Of course, given that the parameters are uncertain, this outcome is not obtainable, but it does provide a benchmark against which we can measure the costs associated with parameter uncertainty.

We now evaluate under parameter uncertainty the optimized simple rules that we computed in the previous section under the assumption that the model parameters were known with certainty. We assume that the central bank's estimates of the natural rates of output are computed using the point estimates of the model parameters, but the actual model parameters and therefore natural rates differ from the values assumed by the policymaker. The central bank is assumed to observe the technology shocks without error since these do not depend on model parameters.

Monetary policy rules optimized assuming no parameter uncertainty are not robust to

reasonable degrees of parameter uncertainty. The third column of the table reports the outcomes under parameter uncertainty for the benchmark estimated model. The welfare loss of about 8 is about 0.5 higher than that under the first-best policies, with most of the increase in the costs associated with output gap fluctuations, and nearly all the remainder from increased fluctuations in wage inflation. A similar result holds for the model estimated only using monetary shocks, for which the results are reported in the final three columns, although the welfare losses in that model are greater. However, in the case of the model estimated using only technology shocks, the model parameters are estimated with very little precision and the policy rule designed for known parameter values performs incredibly poorly under parameter uncertainty.

Most of the increase in the welfare loss under this rule in the benchmark version of the model is due to mismeasurement of the natural rate of output. A strong response to the output gap introduces large policy errors owing to natural rate mismeasurement. The second column reports the results under parameter uncertainty, but assuming that the central bank somehow knew the true values of the natural rate of output, despite not knowing the true parameter values. In this case, the increase in the expected loss relative to the first-best is only 0.1. Thus, in this version of the model at least, absent natural rate mismeasurement, parameter uncertainty would not have much effect on the expected loss under this policy rule. Again, the results are qualitatively similar for the case of the model estimated using monetary shocks, but the welfare losses are huge in the case of the model estimated with technology shocks alone.

We now consider the performance of the policy rules where the coefficients are optimized taking into account parameter uncertainty. We continue to assume no response to wage inflation. The results are shown in Table 5, which follows the same format as Table 4. If the natural rate of output were known, then the optimized policy rule responds much more aggressively to the output gap than under no uncertainty. In this exercise, we have capped the value of ϕ_y at 20 and in all three cases the upper bound is reached. But, such policies would perform very poorly if natural rates were uncertain.

Relative to the case of no uncertainty, the optimized policy rules under parameter uncertainty respond far more aggressively to price inflation, and the response to the output gap, relative to that of inflation, is much smaller. This pattern of policy responses reduces the effects of policy errors associated with mismeasured output gaps. In particular, the

Table 5: Performance of Optimized Policy Rules under Parameter Uncertainty

	Model Estimation								
	Both shocks			Tech shocks			Monetary shocks		
	Opt.	Simple Rule		Opt.	Simple Rule		Opt.	Simple Rule	
Known $\tilde{y}$?	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No
Welfare Losses									
$\mathcal{L}$	7.554	7.566	7.888	10.132	23.426	96.284	19.089	19.758	22.478
$\mathcal{L}_y$	.006	.004	.221	.032	.062	12.318	.306	.085	1.433
$\mathcal{L}_p$	7.258	7.233	7.128	9.136	7.926	5.037	9.563	8.392	8.208
$\mathcal{L}_w$	.292	.329	.539	.964	15.439	78.929	9.220	11.282	11.838
Standard deviations									
x	.023	.006	.085	.066	.042	1.152	.110	.033	.120
π_p	.402	.401	.398	.360	.343	.181	.111	.105	.104
π_w	.017	.018	.022	.086	.117	.289	.139	.153	.156
Policy rule coefficients									
ϕ_p	–	1.57	1.78	–	5.15	20.80	–	1.00	16.07
ϕ_y	–	20.00	.36	–	20.00	.06	–	20.00	20.00

stronger response to inflation works to counteract policy errors resulting from output gap mismeasurement. In two of the cases, including the benchmark version, the response to the output gap is muted relative to the case of no uncertainty. In the case of the model estimated using only monetary shocks, the response to the output gap is larger than under no uncertainty, but the proportional increase in the response to inflation is even larger.

The optimized Taylor-type rules are not completely successful at reducing the welfare costs of fluctuations. In particular, the mean welfare loss in the model estimated using only technology shocks is extremely large. This class of rules is capable of stabilizing price inflation, but does much worse at minimizing the welfare costs associated with fluctuations in the output gap and wage inflation. The problem is that the output gap is a problematic measure of deviations of the real economy from its natural rate and responding overwhelming to price inflation does not correctly balance the joint objectives of stabilizing the output gap, price inflation, and wage inflation.

We now turn to a class of rules that is better suited for this problem, namely, policies that respond explicitly to the wage inflation rate as well as to the output gap and the rate

Table 6: Performance of Optimized Policy Rules under Parameter Uncertainty

	Model Estimation					
	Both shocks		Tech shocks		Monetary shocks	
	Opt.	Simple	Opt.	Simple	Opt.	Simple
Welfare Losses						
$\mathcal{L}$	7.554	7.602	10.132	13.288	19.089	20.824
$\mathcal{L}_y$	.006	.038	.032	.760	.306	1.283
$\mathcal{L}_p$	7.258	7.234	9.136	10.492	9.563	8.524
$\mathcal{L}_w$	.292	.330	.964	2.036	9.220	11.018
Standard deviations						
x	.023	.031	.066	.244	.110	.115
π_p	.402	.401	.360	.409	.111	.106
π_w	.017	.018	.086	.022	.139	.150
Policy rule coefficients						
ϕ_p	–	6.11	–	5.79	–	101.00
ϕ_y	–	.00	–	.01	–	29.92
ϕ_w	–	99.94	–	100.00	–	68.28

of price inflation. Table 6 reports the results. Note that we have restricted the policy rule coefficients to not exceed 100.

Optimized policy rules that respond to wage and price inflation and the output gap are very effective at minimizing the welfare loss under parameter uncertainty. Starting with the benchmark version of the model, the optimized rule responds aggressively to both wage and price inflation and nearly not at all to the output gap. Remarkably, this rule nearly matches the first-best outcome under the assumption that the policymaker knows all the parameter values. In the other two version of the model, the welfare loss under this class of rule is somewhat higher than the ideal first-best, but these rules nonetheless perform noticeably better than optimized standard Taylor rules *assuming that the policymaker knew the true values of the natural rate of output*. Evidently, responding to the wage inflation rate is a highly effective substitute for responding to the output gap in the versions of the model estimated using technology shocks. Indeed, the optimized policies are nearly identical in those two cases, but differ significantly from that in the model estimated using monetary shocks.

Table 7: Performance of Optimized Policy Rules under Parameter Uncertainty

	Optimal Policy	Simple Rule	
Known $\tilde{y}$?	Yes	Yes	No
Welfare Losses			
$\mathcal{L}$	12.259	12.867	28.355
$\mathcal{L}_y$	.115	.183	12.196
$\mathcal{L}_p$	8.652	8.731	13.809
$\mathcal{L}_w$	3.492	3.954	2.351
Standard deviations			
x	.065	.056	.468
π_p	.291	.292	.318
π_w	.081	.080	.040
Policy rule coefficients			
ϕ_p	–	4.26	13.09
ϕ_y	–	100.00	.12
ϕ_w	–	100.00	88.06

5.3 Uncertainty Regarding Estimation Strategy

We now consider the problem of joint uncertainty about parameters and estimation strategy. We implement this by placing equal prior weights on the three estimated models and the associated covariances of parameter estimates. We compute the optimized policy rule that responds to the output gap and the rates of wage and price inflation. The results are reported in Table 7.

Under the optimized rule, joint parameter and estimation strategy uncertainty has a relatively small effect on the mean welfare loss relative to the first-best outcome in the absence of mismeasurement of the output gap. In that case, the optimized rule would respond aggressively to both the output gap and the wage inflation rate.

But, with mismeasurement of the output gap owing to parameter uncertainty, the mean welfare loss under parameter uncertainty is much greater. This single rule is not able to effectively contain the welfare costs across the wide dispersion in parameter estimates across the three models. The optimized rule responds aggressively to both price and wage inflation and responds little to the output gap. The extreme differences between the model estimated using monetary shocks alone and the other two estimated versions of the model are simply

too great for a single policy rule in this class to span.

6 Conclusions

Three crucial factors determine the design of optimal monetary policy: the dynamics of the economy, the natural rate of output, and the weights in the central bank objective function. Traditional analysis of monetary policy under uncertainty has treated these three factors as being independent and studied them separately. But, modern micro-founded models imply that the structural parameters describing preferences and technology jointly determine all three factors and that parameter uncertainty affects all three factors together. This paper has shown in an estimated macroeconomic model that uncertainty about the natural rate of output resulting from parameter uncertainty implies the need for a reorientation of monetary policy from a focus on responding to the output gap to one of responding to wage and price inflation.

References

[1] Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde. "Technology Shocks and Aggregate Fluctuations." Federal Reserve Bank of Chicago, manuscript, June 2002.

[2] Amato, Jeffery, and Thomas Laubach. "Estimation and Control of an Optimization-Based Model with Sticky Prices and Wages." *Journal of Economic Dynamics and Control* 27 (7), May 2003, 1181-1215.

[3] Basu, Susanto, and Miles S. Kimball. "Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption." Manuscript, University of Michigan, October 2002.

[4] Canova, Fabio, and Luca Sala. "Back to Square One: Identification Issues in DSGE Models." European Central Bank, Working Paper 583, January 2006.

[5] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1), February 2005, 1-45.

[6] Clarida, Richard, Jordi Galí, and Mark Gertler. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115 (1), February 2000, 147-180.

[7] Edge, Rochelle, Thomas Laubach, and John Williams. "Learning and Shifts in Long-Run Producitivty Growth." Federal Reserve Board, FEDS paper 2004-21, March 2004.

[8] Finan, Frederico, and Robert Tetlow. "Optimal Control of Large, Forward-Looking Models: Efficient Solutions and Two Examples." Federal Reserve Board, FEDS paper 1999-51, October 1999.

[9] Francis, Neville, and Valerie A. Ramey. "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited." *Journal of Monetary Economics* 52 (8), November 2005, 1379-1399.

[10] Fuhrer, Jeffrey. "Habit Formation in Consumption and Its Implications for Monetary-Policy Models." *American Economic Review* 90 (3), June 2000, 367-390.

[11] Galí, Jordi. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89 (1), March 1999, 249-271.

[12] Galí, Jordi, J. David López-Salido, and Javier Vallés. "Technology Shocks and Monetary Policy: Assessing the Fed's Performance." *Journal of Monetary Economics* 50 (4), May 2003, 723-743.

[13] Giannoni, Marc. "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model." *Macroeconomic Dynamics* 6 (1), February 2002, 111-144.

[14] Laubach, Thomas, and John C. Williams. "The Response of Hours to Technology Shocks: Much Ado About Nothing?" Manuscript, Federal Reserve Bank of San Francisco, November 2006.

[15] Levin, Andrew, Alexei Onatski, John Williams, and Noah Williams. "Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models." *NBER Macroeconomics Annual* 2005, 229-289.

[16] Levin, Andrew, Volker Wieland, and John Williams. "Robustness of Simple Monetary Policy Rules under Model Uncertainty." In John Taylor (ed.), *Monetary Policy Rules*, Chicago, IL: The University of Chicago Press, 1999, 263-299.

[17] Levin, Andrew, and John Williams. "Parameter Uncertainty and the Central Bank's Objective Function." Manuscript, Federal Reserve Bank of San Francisco, October 2005.

[18] Orphanides, Athanasios, and John Williams. "Robust Monetary Policy Rules with Unknown Natural Rates." *Brookings Papers on Economic Activity* 2002 (2), 63-118.

[19] Rotemberg, Julio J. and Michael Woodford. "Real-Business-Cycle Models and the Forecastable Movements in Output, Hours, and Consumption." *American Economic Review* 86 (1), March 1996, 71-89.

[20] Rotemberg, Julio J. and Michael Woodford. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy." *NBER Macroeconomics Annual* 1997, 297-346.

[21] Shapiro, Matthew, and Mark Watson. "Sources of Business Cycle Fluctuations." *NBER Macroeconomics Annual* 1988, 111-148.

[22] Smets, Frank, and Raf Wouters. "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1 (5), September 2003, 1123-1175.

[23] Woodford, Michael. "Optimal Monetary Policy Inertia." NBER Working Paper 7261, August 1999.

A The Linearized Model

The first sub-section of this appendix reports the model's non-linear equations while the second sub-section reports the model's log-linear equations. We limit ourselves throughout in reporting just the equations from the symmetric model. The third sub-section reports the model's natural rate of output.

A.1 First-order conditions

The first-order conditions from the intermediate goods producing firms' cost-minimization problem (equation 8), the labor demand curve and marginal cost function, are:

$$L_{y,t} = \frac{Y_{m,t}}{A_t} \tag{14}$$

$$\frac{MC_t}{P_t} = \frac{W_t}{P_t} \cdot \frac{1}{A_t} \tag{15}$$

The first-order condition from the intermediate goods producing firms profit-maximization problem (equation 9), the aggregate supply curver, is:

$$\begin{aligned} \theta_p \cdot \frac{MC_t}{P_t} \cdot Y_{f,t} = (\theta_p - 1)\left(1+\varsigma_{\theta,p}\right) Y_{f,t} + \chi_p \left(\Pi_{p,t} - \Pi_{p,*}\right) \Pi_{p,t} \cdot \frac{MC_t}{P_t} \cdot Y_{m,t} \\ - \beta E_t \left[\frac{\Lambda_{c,t+1}}{\Lambda_{c,t}} \cdot \chi_p \left(\Pi_{p,t+1} - \Pi_{p,*}\right) \Pi_{p,t+1} \cdot \frac{MC_{t+1}}{P_{t+1}} \cdot Y_{m,t+1} \right] \end{aligned} \tag{16}$$

The first-order conditions from the household's utility-maximization problem (equation 10), the Euler equation and the labor supply curve, are:

$$\frac{\Lambda_{c,t}}{P_t} = \beta R_t E_t \left[\frac{\Lambda_{c,t+1}}{P_{t+1}} \right] \tag{17}$$

$$\begin{aligned} \theta_w \cdot \frac{\Lambda_{l,t}}{\Lambda_{c,t}} \cdot L_{y,t} = (\theta_w - 1)\left(1+\varsigma_{\theta,w}\right) \frac{W_t}{P_t} \cdot L_{y,t} + \chi_w \left(\Pi_{w,t} - \Pi_{w,*}\right) \Pi_{w,t} \cdot \frac{\Lambda_{l,t}}{\Lambda_{c,t}} \cdot L_{u,t} \\ - \beta E_t \left[\frac{\Lambda_{c,t+1}}{\Lambda_{c,t}} \cdot \chi_w \left(\Pi_{w,t+1} - \Pi_{w,*}\right) \Pi_{w,t+1} \cdot \frac{\Lambda_{l,t+1}}{\Lambda_{c,t+1}} \cdot L_{u,t+1} \right] \end{aligned} \tag{18}$$

where

$$\Lambda_{c,t} = \left(C_t - \eta C_{t-1}\right)^{-\sigma} \left(\bar{L} - L_t\right)^{\zeta(1-\sigma)} - E_t \beta \eta \left[\left(C_{t+1} - \eta C_t\right)^{-\sigma} \left(\bar{L} - L_{t+1}\right)^{\zeta(1-\sigma)} \right] \tag{19}$$

$$\Lambda_{l,t} = \zeta \left(C_t - \eta C_{t-1}\right)^{(1-\sigma)} \left(\bar{L} - L_t\right)^{\zeta(1-\sigma)-1}. \tag{20}$$

The model has three market clearing conditions: the labor market clearing condition, the intermediate-goods market clearing condition, and the final-goods market clearing condi-

tion. In the symmetric equilibrium these are given by:

$$L_{u,t} = L_{y,t} + \frac{\chi_w}{2} \left(\Pi_{w,t} - \Pi_{w,t-1}\right)^2 L_{u,t}, \tag{21}$$

$$Y_{f,t} = Y_{m,t} - \frac{\chi_p}{2} \left(\Pi_{p,t} - \Pi_{p,t-1}\right)^2 Y_{m,t}, \tag{22}$$

$$C_t = Y_{f,t} \tag{23}$$

A.2 Log-linearized first-order conditions

The first-order conditions implied by the intermediate goods producing firm's cost minimization problem, given by equations (14) and (15), log-linearize to

$$l_{y,t} = y_{m,t} - a_t \tag{24}$$

$$mc_t = w_t - a_t \tag{25}$$

The first-order conditions implied by the intermediate goods producing firm's profit maximization problem, given by equation (16), log-linearizes to

$$\begin{aligned} \pi_{p,t} &= \beta E_t \pi_{p,t+1} + \frac{(\theta_p - 1)(1 + \varsigma_{\theta,p})}{\Pi_{p,*}^2 \chi_p} \cdot mc_t \\ &= \beta E_t \pi_{p,t+1} + \kappa_p \cdot mc_t \end{aligned} \tag{26}$$

The first-order conditions implied by the household's utility maximization problem, given by equations (17) and (18), log-linearize to:

$$\lambda_{c,t} = r_t - E_t \pi_{p,t+1} + E_t \lambda_{c,t+1} \tag{27}$$

$$\begin{aligned} \pi_{w,t} &= \beta E_t \pi_{w,t+1} + \frac{(\theta_w - 1)(1 + \varsigma_{\theta,l})}{\Pi_{w,*}^2 \chi_w} \cdot (\lambda_{l,t} - \lambda_{c,t} - w_t) \\ &= \beta E_t \pi_{w,t+1} + \kappa_w \cdot (\lambda_{l,t} - \lambda_{c,t} - w_t) \end{aligned} \tag{28}$$

where

$$\begin{aligned} \lambda_{c,t} = \frac{1}{1-\eta\beta} \Bigg(& \frac{-\sigma}{1-\eta} (c_t - \eta c_{t-1}) - \zeta(1-\sigma) \frac{L_{u,*}}{\bar{L} - L_{u,*}} \cdot l_{u,t} \Bigg) \\ & - \frac{\eta\beta}{1-\eta\beta} \Bigg(\frac{-\sigma}{1-\eta} (E_t c_{t+1} - \eta c_t) - \zeta(1-\sigma) \frac{L_{u,*}}{\bar{L} - L_{u,*}} \cdot E_t l_{u,t+1} \Bigg) \end{aligned} \tag{29}$$

$$\lambda_{l,t} = \frac{1-\sigma}{1-\eta} (c_t - \eta c_{t-1}) + (1 - \zeta(1-\sigma)) \frac{L_{u,*}}{\bar{L} - L_{u,*}} \cdot l_{u,t}. \tag{30}$$

The market clearing conditions, equations (21), (22), and (23), log-linearize to:

$$l_{u,t} = l_{y,t} = l_t \tag{31}$$

$$y_{f,t} = y_{m,t} = y_t \tag{32}$$

$$c_t = y_{f,t} \tag{33}$$

Three more equations remain in our model: (i) the process for the shocks A_t, defined in equation (4), which log-linearizes to:

$$a_t = a_{t-1} + \epsilon_t, \tag{34}$$

(ii) the monetary policy process, which was already given in log-linearized form in equation (11), and (iii) an identity between price and wage inflation and real wages:

$$\frac{W_t}{P_t} \cdot \frac{P_{t-1}}{W_{t-1}} = \frac{W_t}{W_{t-1}} \cdot \frac{P_{t-1}}{P_t} = \frac{\Pi_{w,t}}{\Pi_{p,t}}$$

which log-linearized to:

$$w_t - w_{t-1} = \pi_{w,t} - \pi_{p,t}$$

Before concluding this section we note the following about the steady-state solution to the model. We know from equations (15), (16), (18), (19), and (20) that in the steady state:

$$1 = \frac{MC_*}{P_*} = \frac{W_*}{P_*} = \frac{\Lambda_{l,*}}{\Lambda_{c,*}} = \frac{1-\eta}{1-\eta\beta} \cdot \frac{\zeta L_{u,*}}{\bar{L} - L_{u,*}}$$

Since $L_* = L_{u,*} = L_{y,*}$ we can re-write this as:

$$\frac{\zeta L_*}{\bar{L} - L_*} = \frac{1-\eta\beta}{1-\eta}. \tag{35}$$

This means that we can re-write equations (29) and (30) as:

$$\begin{aligned}
\lambda_{c,t} = \frac{1}{1-\eta\beta}\left(\frac{-\sigma}{1-\eta}\left(c_t - \eta c_{t-1}\right) - (1-\sigma)\frac{1-\eta\beta}{1-\eta} \cdot l_{u,t}\right) \\
+ \frac{\eta\beta}{1-\eta\beta}\left(\frac{-\sigma}{1-\eta}(E_t c_{t+1} - \eta c_t) - (1-\sigma)\frac{1-\eta\beta}{1-\eta} \cdot E_t l_{u,t+1}\right) \\
\lambda_{l,t} = \frac{1-\sigma}{1-\eta}\left(c_t - \eta c_{t-1}\right) + \frac{1-\zeta\left(1-\sigma\right)}{\zeta} \cdot \frac{1-\eta\beta}{1-\eta} \cdot l_{u,t}.
\end{aligned}$$

Since $l_{u,t} = l_{y,t} = y_{m,t} - a_t = y_t - a_t$ and $c_t = y_{f,t} = y_t$

$$\begin{aligned}
\lambda_{c,t} = \frac{1}{1-\eta\beta}\left(\frac{-\sigma}{1-\eta}\left(y_t - \eta y_{t-1}\right) - (1-\sigma)\frac{1-\eta\beta}{1-\eta}(y_t - a_t)\right) \\
+ \frac{\eta\beta}{1-\eta\beta}\left(\frac{-\sigma}{1-\eta}(E_t y_{t+1} - \eta y_t) - (1-\sigma)\frac{1-\eta\beta}{1-\eta}(E_t y_{t+1} - E_t a_{t+1})\right) \qquad (36) \\
\lambda_{l,t} = \frac{1-\sigma}{1-\eta}\left(y_t - \eta y_{t-1}\right) + \frac{1-\zeta\left(1-\sigma\right)}{\zeta} \cdot \frac{1-\eta\beta}{1-\eta}(y_t - a_t). \qquad (37)
\end{aligned}$$

Equations (36) and (37) will be used in deriving the natural rate.

A.3 Natural rates in the log-linear model

The natural rate of output, i.e. the level of output in the equilibrium with perfectly flexible prices and wages, is determined by the condition that the marginal rate of substitution between consumption and leisure be (up to constants) equal to the marginal product of labor at all dates t. In log terms, $\tilde{y}_t$ is determined implicitly by the equation

$$mrs_t = mpl_t \tag{38}$$

where $mrs_t = \lambda_{l,t} - \lambda_{c,t}$ and, from the production function (3), $mpl_t = a_t$. Substituting from (36) and (37) for $\lambda_{c,t}$ and $\lambda_{l,t}$ yields the following expression for $\tilde{y}_t$:

$$[\delta_1 - \delta_2(L + \beta L^{-1})]\tilde{y}_t = [\delta_3 - \delta_4 L^{-1}]a_t \tag{39}$$

where L denotes the lag operator, $L^{-1}x_t \equiv E_t x_{t+1}$, and

$$\begin{aligned}
\delta_1 &= \frac{2(1-\sigma)}{1-\eta} + \frac{1-\zeta(1-\sigma)}{\zeta}\frac{1-\beta\eta}{1-\eta} + \frac{\sigma(1+\beta\eta^2)}{(1-\beta\eta)(1-\eta)} \\
\delta_2 &= \eta\left[\frac{1-\sigma}{1-\eta} + \frac{\sigma}{(1-\beta\eta)(1-\eta)}\right] \\
\delta_3 &= 1 + \frac{1-\zeta(1-\sigma)}{\zeta}\frac{1-\beta\eta}{1-\eta} + \frac{1-\sigma}{1-\eta}, \text{ and} \\
\delta_4 &= \frac{\beta\eta(1-\sigma)}{1-\eta}
\end{aligned}$$

The natural rate of interest, denoted $\tilde{r}_t$, is the real rate $r_t - E_t\pi_{t+1}$ prevailing in the equilibrium with perfectly flexible prices and wages. Letting $\tilde{\lambda}_{c,t}$ denote the expression (36) with $\tilde{y}_t$ substituted for y_t, the Euler equation (27) in this equilibrium can be expressed as

$$\tilde{r}_t = \tilde{\lambda}_{c,t} - E_t\tilde{\lambda}_{c,t+1} \tag{40}$$

But, in this equilibrium, $\tilde{\lambda}_{c,t} = \tilde{\lambda}_{l,t} - a_t$. Substituting for $\tilde{\lambda}_{c,t}$ we obtain

$$\begin{aligned}
\tilde{r}_t &= [\delta_6 + \delta_7 - \delta_6 L - \delta_7 L^{-1}]\tilde{y}_t + \delta_8[L^{-1} - 1]a_t \\
&= \delta_6\Delta\tilde{y}_t - \delta_7 E_t\Delta\tilde{y}_{t+1} + \delta_8 E_t\Delta a_{t+1}
\end{aligned} \tag{41}$$

with

$$\begin{aligned}
\delta_6 &= \frac{\eta(1-\sigma)}{1-\eta} \\
\delta_7 &= \frac{1-\sigma}{1-\eta} + \frac{1-\zeta(1-\sigma)}{\zeta}\frac{1-\beta\eta}{1-\eta} \\
\delta_8 &= 1 + \frac{1-\zeta(1-\sigma)}{\zeta}\frac{1-\beta\eta}{1-\eta}
\end{aligned}$$

B VAR Identification and Structural Shocks

As discussed in section 3, we are interested in matching the VAR's impulse responses to two of the structural shocks of our model, a permanent shock to the level of technology, and a transitory shock to the funds rate. To identify these shocks, we use one long-run and one short-run identifying restriction. The short-run identifying restriction is the usual one, that the last variable in the VAR (the funds rate) is Wold-causal for the preceding variables. The structural form of the VAR is given by

$$A_0 Y_t = constant + A(L)Y_{t-1} + \varepsilon_t, \tag{42}$$

where Y_t is defined in (12). The short-run assumption implies that the last column of the contemporaneous multiplier matrix A_0 has all zeros above the main diagonal. The fifth element of ε_t is identified as the funds rate shock $\epsilon_{r,t}$ in (11). The long-run identifying restriction is the one proposed by Galí (1999) and further explored in Francis and Ramey (2005) and Altig *et al.* (2002), that permanent shocks to technology are the only shocks to have a permanent effect on labor productivity. Using this assumption, we identify the first element of ε_t as the technology shock ϵ_t in (4). This implies that the first row of the matrix of long-run (cumulative) effects of ε_t on Y_t, $(I - A(1))^{-1}A_0^{-1}$, consists of zeros except for the first element.

In order to estimate the VAR in structural form, we need a further set of assumptions to just-identify the elements of A_0. We follow Altig *et al.* (2002) by assuming that the submatrix consisting of columns 2-4 and rows 2-4 of A_0 is lower triangular. This assumption is without loss of generality as we do not attach any structural interpretation to elements 2 through 4 of ε_t. With these assumptions, we estimate the first equation of the structural VAR imposing the long-run restrictions in the manner of Shapiro and Watson (1988) by including contemporaneous and lagged variables of elements 2 through 4 of Y_t in first-differenced form. To control for simultaneity, we estimate the equation by 2SLS, using a constant and $Y_{t-1}, \ldots, Y_{t-4}$ as first-stage regressors for elements 2 through 4 of Y_t. We then sequentially estimate equations 2 through 4 by IV, using the residuals from the previous regressions as instruments for contemporaneous variables. Equation 5 can be estimated by OLS by virtue of our short-run identifying assumption.

We modify this identification strategy in one respect. Because, in contrast to Altig *et al.*, our VAR includes hours per capita in first differences, we would like to assure that the *long-*

run response of hours to a technology shock is zero, consistent with the observation that hours worked have remained broadly unchanged despite the secular trend in real wages. When this second long-run restriction is not imposed, the IRF of Δl usually does not integrate to zero. We therefore reorder our vector of endogenous variables to include Δl as the second variable, and apply the Shapiro-Watson method to the first two equations. This leaves the interpretation of the first element of ε_t unchanged, but the second element is now the only shock that permanently affects hours per capita. Contrary to the findings reported by Francis and Ramey (2005), Laubach and Williams (2006) find that imposing this second long-run restriction can have a substantial effect on the response of hours to a technology shock.

C Deriving the Welfare Criterion

To derive the welfare criterion we first take a second-order approximation to the within-period utility function

$$\begin{aligned}
&\frac{1}{1-\sigma}\cdot\frac{\left(C_t-\eta C_{t-1}\right)^{1-\sigma}\left(\bar{L}-L_{u,t}\right)^{\zeta(1-\sigma)}}{\left(C_*-\eta C_*\right)^{1-\sigma}\left(\bar{L}-L_{u,*}\right)^{\zeta(1-\sigma)}}\\
&\quad=\text{T.I.P.}\\
&\quad+\frac{1}{1-\eta}\left(c_t+\frac{1}{2}\cdot c_t^2\right)+\frac{-\eta}{1-\eta}\left(c_{t-1}+\frac{1}{2}\cdot c_{t-1}^2\right)-\frac{\zeta L_{u,*}}{\bar{L}-L_{u,*}}\left(l_{u,t}+\frac{1}{2}\cdot l_{u,t}^2\right)\\
&\quad+\frac{1}{2}(-\sigma)\frac{1}{1-\eta}\cdot\frac{1}{1-\eta}\left(c_t+\frac{1}{2}\cdot c_t^2\right)\left(c_t+\frac{1}{2}\cdot c_t^2\right)\\
&\quad+\frac{1}{2}(-\sigma)\frac{-\eta}{1-\eta}\cdot\frac{-\eta}{1-\eta}\left(c_{t-1}+\frac{1}{2}\cdot c_{t-1}^2\right)\left(c_{t-1}+\frac{1}{2}\cdot c_{t-1}^2\right)\\
&\quad-\frac{1}{2}\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\frac{\zeta L_{u,*}}{\bar{L}-L_{u,*}}\cdot\frac{\zeta L_{u,*}}{\bar{L}-L_{u,*}}\left(l_{u,t}+\frac{1}{2}\cdot l_{u,t}^2\right)\left(l_{u,t}+\frac{1}{2}\cdot l_{u,t}^2\right)\\
&\quad+(-\sigma)\frac{1}{1-\eta}\cdot\frac{-\eta}{1-\eta}\left(c_t+\frac{1}{2}\cdot c_t^2\right)\left(c_{t-1}+\frac{1}{2}\cdot c_{t-1}^2\right)\\
&\quad-(1-\sigma)\frac{1}{1-\eta}\cdot\frac{\zeta L_{u,*}}{\bar{L}-L_{u,*}}\left(c_t+\frac{1}{2}\cdot c_t^2\right)\left(l_{u,t}+\frac{1}{2}\cdot l_{u,t}^2\right)\\
&\quad-(1-\sigma)\frac{-\eta}{1-\eta}\cdot\frac{\zeta L_{u,*}}{\bar{L}-L_{u,*}}\left(c_{t-1}+\frac{1}{2}\cdot c_{t-1}^2\right)\left(l_{u,t}+\frac{1}{2}\cdot l_{u,t}^2\right) \qquad (43)
\end{aligned}$$

We also make use of the quadratic approximations to the labor demand curve (equation 14) and the market clearing conditions (equations 22, 21, and 23) which are given by:

$$\left(l_{y,t}+\frac{1}{2}\cdot l_{y,t}^2\right)=\left(y_{m,t}+\frac{1}{2}\cdot y_{m,t}^2\right)-\left(a_t+\frac{1}{2}\cdot a_t^2\right)-y_{m,t}a_t$$

$$\left(l_{y,t}+\frac{1}{2}\cdot l_{y,t}^2\right)=\left(l_{u,t}+\frac{1}{2}\cdot l_{u,t}^2\right)-\frac{1}{2}\cdot\chi_w\cdot\Pi_{w,*}^2\pi_{w,t}^2$$
$$\left(y_{f,t}+\frac{1}{2}\cdot y_{f,t}^2\right)=\left(y_{m,t}+\frac{1}{2}\cdot y_{m,t}^2\right)-\frac{1}{2}\cdot\chi_p\cdot\Pi_{p,*}^2\pi_{p,t}^2$$
$$\left(c_t+\frac{1}{2}\cdot c_t^2\right)=\left(y_{f,t}+\frac{1}{2}\cdot y_{f,t}^2\right)$$

Substituting the above equations into the equation (43) and making a number of substitutions yields:

$$\begin{aligned}
&\frac{1}{1-\sigma}\cdot\frac{\left(C_t-\eta C_{t-1}\right)^{1-\sigma}\left(\bar{L}-L_{u,t}\right)^{\zeta(1-\sigma)}}{\left(C_*-\eta C_*\right)^{1-\sigma}\left(\bar{L}-L_{u,*}\right)^{\zeta(1-\sigma)}}\\
&\quad=\text{T.I.P.}\\
&\quad-\frac{1}{2}\cdot\sigma\left(\frac{1}{1-\eta}\right)^2\left(\left(y_t-\widetilde{y}_t\right)-\eta\left(y_{t-1}-\widetilde{y}_{t-1}\right)\right)^2\\
&\quad-\frac{1}{2}\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\left(\frac{1-\beta\eta}{1-\eta}\right)^2\left(y_t-\widetilde{y}_t\right)^2\\
&\quad-(1-\sigma)\cdot\frac{1-\beta\eta}{1-\eta}\cdot\frac{1}{1-\eta}\cdot\left(y_t-\widetilde{y}_t\right)\left(\left(y_t-\widetilde{y}_t\right)-\eta\left(y_{t-1}-\widetilde{y}_{t-1}\right)\right)\\
&\quad-\frac{1}{2}\cdot\frac{1-\beta\eta}{1-\eta}\left\{\chi_p\Pi_{p,*}^2\pi_{p,t}^2+\ \chi_w\Pi_{w,*}^2\pi_{w,t}^2\right\}.
\end{aligned}$$

Letting $x_t\equiv y_t-\widetilde{y}_t$, and taking account our expressions for $\kappa_{r,p}$ and $\kappa_{r,w}$ this can also be written as:

$$\begin{aligned}
&\frac{1}{1-\sigma}\cdot\frac{\left(C_t-\eta C_{t-1}\right)^{1-\sigma}\left(\bar{L}-L_{u,t}\right)^{\zeta(1-\sigma)}}{\left(C_*-\eta C_*\right)^{1-\sigma}\left(\bar{L}-L_{u,*}\right)^{\zeta(1-\sigma)}}\\
&\quad=\text{T.I.P.}\\
&\quad-\frac{1}{2}\cdot\sigma\left(\frac{1}{1-\eta}\right)^2\left(x_t-\eta\,x_{t-1}\right)^2\\
&\quad-\frac{1}{2}\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\left(\frac{1-\beta\eta}{1-\eta}\right)^2x_t^2\\
&\quad-(1-\sigma)\cdot\frac{1-\beta\eta}{1-\eta}\cdot\frac{1}{1-\eta}\cdot x_t\left(x_t-\eta\,x_{t-1}\right)\\
&\quad-\frac{1}{2}\cdot\frac{1-\beta\eta}{1-\eta}\left\{\frac{\theta_p}{\kappa_{r,p}}\pi_{p,t}^2+\ \frac{\theta_w}{\kappa_{r,w}}\pi_{w,t}^2\right\},
\end{aligned}$$

which is the equation given in section 4 of the paper.

D The Model's Calvo Counterpart

As noted in section 2.3 of the paper we preferred to use the quadratic adjustment cost approach to modeling sticky-prices and sticky-wages in preference to alternatives that imply

heterogeneity among agents. Naturally, this raises the question of whether this modelling choice has serious implications for our results.

In this appendix we present our model's Calvo counterpart. Clearly the principal differences are that we no longer have the price and wage adjustment cost terms in equations 2 and 6 (so that these equations become simply $Y_{f,t}(j) = Y_{m,t}(j)$ and $L_{u,t}(i) = L_{y,t}(i)$); instead, however, the model contains calvo reset probabilities in the firms' profit-maximization problem and the households' utility maximization problems. All of the parts of the model—that is, the firms' cost-minimization problem and the process describing monetary policy—remain the same.

Under Calvo pricing-setting the firm's profit-maximization problem is:

$$\begin{aligned}\max_{\{P_t(j)\}_{t=0}^{\infty}} & \{(1+\varsigma_{\theta,p})Y_{f,0}(j)P_0(j) - Y_{m,0}(j)MC_0(j)\} \\ & +E_0\sum_{t=1}^{\infty}(\beta\alpha_p)^t\frac{\Lambda_{c,t}}{P_t}\left\{(1+\varsigma_{\theta,p})Y_{f,t}(j)P_t(j)\Pi_{p,*}^t - Y_{m,t}(j)MC_t(j)\right\} \\ & \text{subject to } Y_{f,t}(j)=Y_{m,t}(j) \text{ and } Y_{f,t}(j)=(P_t(j)/P_t)^{-\theta_p}Y_{f,t}. \end{aligned} \tag{44}$$

Under Calvo wage-setting the household's utility-maximization problem is:

$$\max_{\{C_t(i),W_t(i)\}_{t=0}^{\infty}} E_0\frac{1}{1-\sigma}\sum_{t=0}^{\infty}\beta^t\left[(C_t(i)-\eta C_{t-1}(i))(\bar{L}-L_{u,t}(i))^{\zeta}\right]^{1-\sigma}$$

subject to

$$E_0\left[\beta\frac{\Lambda_{c,1}/P_{c,1}}{\Lambda_{c,0}/P_{c,0}}B_1(i)\right] = B_0(i) + (1+\varsigma_{\theta,w})L_{y,0}(i)W_0(i) + Profits_0(i) - P_0C_0(i),$$

$$E_t\left[\beta\frac{\Lambda_{c,t+1}/P_{c,t+1}}{\Lambda_{c,t}/P_{c,t}}B_{t+1}(i)\right] = B_t(i) + (1+\varsigma_{\theta,w})L_{y,t}(i)W_t(i)\alpha_w^t + Profits_t(i) - P_tC_t(i),$$

$$\text{for } t = 1, \cdots, \infty.$$

$$L_{y,t}(i) = L_{u,t}(i) \text{ and } L_{y,t}(i) = (W_t(i)/W_t)^{-\theta_w}\int_0^1 L_{y,t}(i,j)dj. \tag{45}$$

All of the equations in the log-linearized model other than the price and wage Phillips curves are unaltered by our assumption for price- and wage- setting. Equations (26) and (28) are however replaced by:

$$\pi_{p,t} = \beta E_t\pi_{p,t+1} + \frac{(1-\alpha_p)(1-\beta\alpha_p)}{\alpha_p}\cdot mc_t \tag{46}$$

$$\begin{aligned}\pi_{w,t} = & \beta E_t\pi_{w,t+1} \\ & + \frac{(1-\alpha_w)(1-\beta\alpha_w)}{\alpha_w}\left(1+\theta_w\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\frac{1-\beta\eta}{1-\eta}\right)^{-1}(\lambda_{l,t}-\lambda_{c,t}-w_t)\end{aligned} \tag{47}$$

In estimating our model, we could have set:

$$\kappa_{c,p} = \frac{(1-\alpha_p)(1-\beta\alpha_p)}{\alpha_p} \text{ and } \kappa_{c,w} = \frac{(1-\alpha_w)(1-\beta\alpha_w)}{\alpha_w}\left(1+\theta_w \cdot \frac{1-\zeta(1-\sigma)}{\zeta}\cdot\frac{1-\beta\eta}{1-\eta}\right)^{-1}$$

and included:

$$\pi_{p,t} = \beta E_t \pi_{p,t+1} + \kappa_{c,p} \cdot mc_t \tag{48}$$

$$\pi_{w,t} = \beta E_t \pi_{w,t+1} + \kappa_{c,w} \cdot (\lambda_{l,t} - \lambda_{c,t} - w_t) \tag{49}$$

in our estimated model. These would have been identical to equations (26) and (28) which we did estimate.

Under Calvo pricing the loss function is:

$$\begin{aligned}
&\frac{1}{1-\sigma}\cdot\frac{\left(C_t-\eta C_{t-1}\right)^{1-\sigma}\left(\bar{L}-L_{u,t}\right)^{\zeta(1-\sigma)}}{\left(C_*-\eta C_*\right)^{1-\sigma}\left(\bar{L}-L_{u,*}\right)^{\zeta(1-\sigma)}}\\
&= \text{T.I.P.}\\
&-\frac{1}{2}\cdot\sigma\left(\frac{1}{1-\eta}\right)^2\left(\left(y_t-\widetilde{y}_t\right)-\eta\left(y_{t-1}-\widetilde{y}_{t-1}\right)\right)^2\\
&-\frac{1}{2}\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\left(\frac{1-\beta\eta}{1-\eta}\right)^2\left(y_t-\widetilde{y}_t\right)^2\\
&-(1-\sigma)\cdot\frac{1-\beta\eta}{1-\eta}\cdot\frac{1}{1-\eta}\cdot\left(y_t-\widetilde{y}_t\right)\left(\left(y_t-\widetilde{y}_t\right)-\eta\left(y_{t-1}-\widetilde{y}_{t-1}\right)\right)\\
&-\frac{1}{2}\cdot\frac{1-\beta\eta}{1-\eta}\cdot\theta_p\cdot\frac{\alpha_p}{(1-\alpha_p)(1-\beta\alpha_p)}\cdot\pi_{p,t}^2\\
&-\frac{1}{2}\cdot\frac{1-\beta\eta}{1-\eta}\cdot\theta_w\cdot\frac{\alpha_w}{(1-\alpha_w)(1-\beta\alpha_w)}\cdot\left(1+\theta_w\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\frac{1-\beta\eta}{1-\eta}\right)^2\cdot\pi_{w,t}^2,
\end{aligned}$$

This can also be written as:

$$\begin{aligned}
&\frac{1}{1-\sigma}\cdot\frac{\left(C_t-\eta C_{t-1}\right)^{1-\sigma}\left(\bar{L}-L_{u,t}\right)^{\zeta(1-\sigma)}}{\left(C_*-\eta C_*\right)^{1-\sigma}\left(\bar{L}-L_{u,*}\right)^{\zeta(1-\sigma)}}\\
&= \text{T.I.P.}\\
&-\frac{1}{2}\cdot\sigma\left(\frac{1}{1-\eta}\right)^2\left(x_t-\eta x_{t-1}\right)^2-\frac{1}{2}\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\left(\frac{1-\beta\eta}{1-\eta}\right)^2 x_t^2\\
&-(1-\sigma)\cdot\frac{1-\beta\eta}{1-\eta}\cdot\frac{1}{1-\eta}\cdot x_t\left(x_t-\eta x_{t-1}\right)\\
&-\frac{1}{2}\cdot\frac{1-\beta\eta}{1-\eta}\cdot\frac{\theta_p}{\kappa_{c,p}}\cdot\pi_{p,t}^2-\frac{1}{2}\cdot\frac{1-\beta\eta}{1-\eta}\cdot\frac{\theta_w}{\kappa_{c,w}}\cdot\underbrace{\left(1+\theta_w\cdot\frac{1-\zeta(1-\sigma)}{\zeta}\cdot\frac{1-\beta\eta}{1-\eta}\right)}_{*}\cdot\pi_{w,t}^2,
\end{aligned}$$

where the "$*$" term is only term that differs from what appears in the loss funciton under our quadratic adjustment costs assumption.