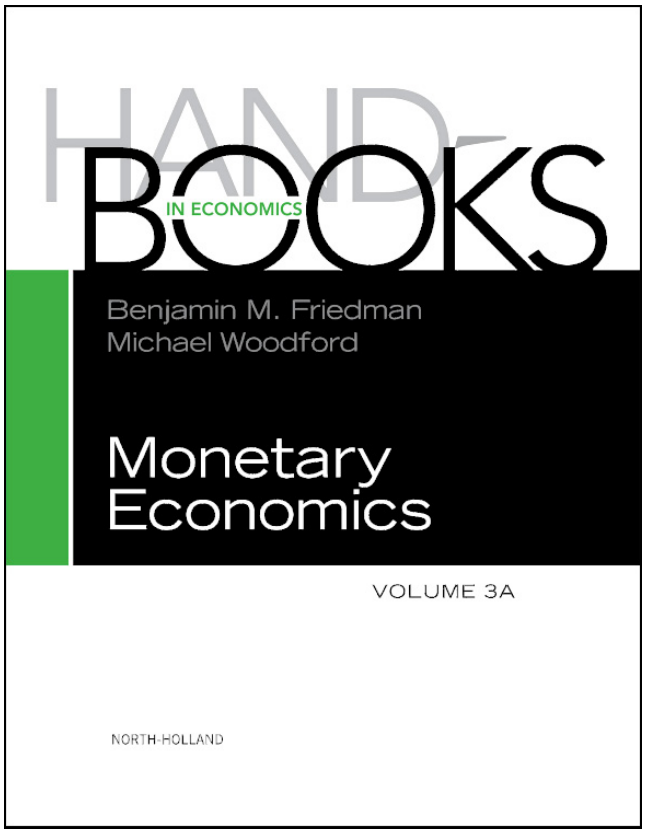

From: Lawrence J. Christiano, Mathias Trabandt, and Karl Walentin, DSGE Models for Monetary Policy Analysis, In Benjamin M. Friedman, and Michael Woodford, editors: Handbook of Monetary Economics, Vol. 3A, The Netherlands: North-Holland, 2011, pp. 285-367.
ISBN: 978-0-444-53238-1

CHAPTER 7

DSGE Models for Monetary Policy Analysis☆

Lawrence J. Christiano,* Mathias Trabandt, and Karl Walentin†**

*Department of Economics, Northwestern University
**European Central Bank, Germany and Sveriges Riksbank, Sweden
†Research Division, Sveriges Riksbank, Sweden

Contents

☆ We are grateful for advice from Michael Woodford and for comments from Volker Wieland. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the European Central Bank or of Sveriges Riksbank. We are grateful for assistance from Daisuke Ikeda and Matthias Kehrig.

Handbook of Monetary Economics, Volume 3A
ISSN 0169-7218, DOI: 10.1016/S0169-7218(11)03007-3

Abstract

Monetary DSGE models are widely used because they fit the data well and they can be used to address important monetary policy questions. We provide a selective review of these developments. Policy analysis with DSGE models requires using data to assign numerical values to model parameters. The chapter describes and implements Bayesian moment matching and impulse response matching procedures for this purpose.
JEL Classification: E2, E3, E5, J6

Keywords

Frisch Labor Supply Elasticity
HP Filter
Impulse Response Function
Limited Information Bayesian Estimation
Materials Input for Production
New Keynesian DSGE Models
Output Gap
Potential Output
Taylor Principle
Unemployment
Vector Autoregression
Working Capital Channel

1. INTRODUCTION

There has been enormous progress in recent years in the development of dynamic, stochastic general equilibrium (DSGE) models for the purpose of monetary policy analysis. These models have been shown to fit aggregate data well by conventional econometric measures. For example, they have been shown to do as well or better than simple atheoretical statistical models at forecasting outside the sample of data on which they were estimated. In part because of these successes, a consensus has formed around a particular model structure, the New Keynesian model.

Our objective is to present a selective review of these developments. We present several examples to illustrate the kind of policy questions the models can be used to address. We also convey a sense of how well the models fit the data. In all cases, our discussion takes place in the simplest version of the model required to make our point. As a result, we do not develop one single model. Instead, we work with several models.

We begin by presenting a detailed derivation of a version of the standard New Keynesian model with price-setting frictions and no capital or other complications. We then use versions of this simple model to address several important policy issues. For example, the past few decades have witnessed the emergence of a consensus that monetary policy ought to respond aggressively to changes in actual or expected inflation. This prescription for monetary policy is known as the "Taylor principle." The standard version of the simple model is used to articulate why this prescription is a good one. However, alternative versions of the model can be used to identify potential pitfalls for the Taylor principle. In particular, a policy-induced rise in the nominal interest rate may destabilize the economy by perversely giving a direct boost to inflation. This can happen if the standard model is modified to incorporate a so-called working capital channel, which corresponds to the assumption that firms must borrow to finance their variable inputs.

We then turn to the much discussed issue of the interaction between monetary policy and volatility in asset prices and other aggregate economic variables. We explain how vigorous application of the Taylor principle could inadvertently trigger an inefficient boom in output and asset prices.

Finally, we discuss the use of DSGE models for addressing a key policy question: How big is the gap between the level of economic activity and the best level that is achievable by policy? An estimate of the output gap not only provides an indication about how efficiently resources are being used, but in the New Keynesian framework, the output gap is also a signal of inflation pressure. Informally, the unemployment rate is thought to provide a direct observation on the efficiency of resource allocation. For example, a large increase in the number of people reporting to be "ready and willing to work" but not employed suggests, at least at a casual level, that resources are being wasted and that the output gap is negative. DSGE models can be used to formalize and assess these informal hunches. We do this by introducing unemployment into the standard New Keynesian model along the lines recently proposed in Christiano, Trabandt, and Walentin (2010a; CTW). We use the model to describe circumstances in which we can expect the unemployment rate to provide useful information about the output gap. We also report evidence suggesting that these conditions may be satisfied in the U.S. data.

Although the creators of the Hodrick and Prescott (1997; HP) filter never intended it to be used to estimate the New Keynesian output gap concept, it is often used for this purpose. We show that whether the HP filter is a good estimator of the gap depends sensitively on the details of the underlying model economy. This discussion involves a careful review of the intuition of how the New Keynesian model responds to shocks. Interestingly, a New Keynesian model fit to U.S. data suggests the conditions are satisfied for the HP filter to be a good estimator of the output gap. In our

discussion, we explain that there are several caveats that must be taken into account before concluding that the HP filter is a good estimator of the output gap.

Policy analysis with DSGE models, even the simple analyses summarized earlier, require assigning values to model parameters. In recent years, the Bayesian approach to econometrics has taken over as the dominant one for this purpose. In conventional applications, the Bayesian approach is a so-called full information procedure because the analyst specifies the joint likelihood of the available observations in complete detail. As a result, many of the limited information tools in macroeconomists' econometric toolbox have been deemphasized in recent times. These tools include methods that match model and data second moments and that match model and empirical impulse response functions. Following the work of Chernozhukov and Hong (2003), Kim (2002), Kwan (1999) and others, we show how the Bayesian approach can be applied in limited information contexts. We apply a Bayesian moment matching approach in Section 3.3.3 and a Bayesian impulse response function matching approach in Section 5.2.

The new monetary DSGE models are of interest not just because they represent laboratories for the analysis of important monetary policy questions. They are also of interest because they appear to resolve a classic empirical puzzle about the effects of monetary policy. It has long been thought that it is virtually impossible to explain the very slow response of inflation to a monetary disturbance without appealing to completely implausible assumptions about price frictions (see, e.g., Mankiw, 2000). However, it turns out that modern DSGE models do provide an account of the inertia in inflation and the strong response of real variables to monetary policy disturbances, without appealing to seemingly implausible parameter values. Moreover, the models simultaneously explain the dynamic response of the economy to other shocks. We review these important findings. We explain in detail the contribution of each feature of the consensus medium-sized New Keynesian model in achieving this result. This discussion closely follows the analyses in Christiano, Eichenbaum, and Evans (2005; CEE) and Altig, Christiano, Eichenbaum, and Lindé (2005; ACEL).

There is an econometric technique that is particularly well-suited to the shock-based analysis described in the previous paragraph. It is the one that matches impulse response functions estimated by vector autoregressions (VARs) with the corresponding objects in a model. Using U.S. macroeconomic data, we show how the parameters of the consensus DSGE model are estimated by this impulse response matching procedure. The advantage of this econometric approach is transparency and focus. The transparency reflects that the estimation strategy has a simple graphical representation, involving objects — impulse response functions — about which economists have strong intuition. The advantage of focus comes from the possibility of studying the empirical properties of a model without having to specify a full set of shocks. As noted previously, we show how to implement the impulse response matching strategy using Bayesian methods. In particular, we are able to implement all the machinery of priors and posteriors, as well as the marginal likelihood as a measure of model fit in our impulse response function matching exercise.

This chapter is organized as follows. Section 2 describes the simple New Keynesian model without capital. The following section reviews some policy implications of that model. The medium-sized version of the model, designed to econometrically address a rich set of macroeconomic data, is described in Section 4. Section 5 reviews our Bayesian impulse response matching strategy. Section 6 reviews the results, and conclusions are offered in Section 7. Many algebraic derivations are relegated to a separate technical appendix.[1]

2. SIMPLE MODEL

This section analyzes versions of the standard Calvo-sticky price New Keynesian model without capital. In practice, the analysis of the standard New Keynesian model often begins with the familiar three equations: the linearized "Phillips curve," "IS curve," and monetary policy rule. We cannot simply begin with these three equations here because we also study departures from the standard model. For this reason, we must derive the equilibrium conditions from their foundations.

The version of the New Keynesian model studied in this section is the one considered in Clarida, Gali, and Gertler (1999) and Woodford (2003), modified in two ways. First, we introduce the working capital channel emphasized by CEE and Barth and Ramey (2002).[2] The working capital channel results from the assumption that firms' variable inputs must be financed by short-term loans. With this assumption, changes in the interest rate affect the economy by changing firms' variable production costs, in addition to operating through the usual spending mechanism. There are several reasons to take the working capital channel seriously. Using U.S. Flow of Funds data, Barth and Ramey (2002) argued that a substantial fraction of firms' variable input costs are borrowed in advance. Christiano, Eichenbaum, and Evans (1996) provided VAR evidence suggesting the presence of a working capital channel. Chowdhury, Hoffmann, and Schabert (2006) and Ravenna and Walsh (2006) provided additional evidence supporting the working capital channel, based on instrumental variables estimates of a suitably modified Phillips curve. Finally, Section 4 shows that incorporating the working capital channel helps to explain the "price puzzle" in the VAR literature and provides a response to Ball's (1994) "dis-inflationary boom" critique of sticky price models.

We explore a second modification to the classic New Keynesian model by incorporating the assumption about materials inputs proposed in Basu (1995). Basu argued that a large part — as much as half — of a firm's output is used as inputs by other firms. The working capital channel introduces the interest rate into costs while the materials assumption makes those costs big. In the next section we show that these two factors have potentially far-reaching consequences for monetary policy.

[1] The technical appendix can be found at http://www.faculty.econ.northwestern.edu/faculty/christiano/research/Handbook/technical_appendix.pdf.

[2] The first monetary DSGE model we are aware of that incorporates a working capital channel is Fuerst (1992). Other early examples include Christiano (1991) and Christiano and Eichenbaum (1992b).

This section is organized as follows. We begin in subsection 2.1 by describing the private sector of the economy, and deriving equilibrium conditions associated with optimization and market clearing. In subsection 2.2, we specify the monetary policy rule and define the Taylor rule equilibrium. Subsection 2.3 discusses the interpretation of a key parameter in our utility function. The parameter controls the elasticity with which the labor input in our model economy adjusts in response to a change in the real wage. Traditionally, this parameter has been viewed as being restricted by microeconomic evidence on the Frisch labor supply elasticity. We summarize recent thinking stimulated by the seminal work of Rogerson (1988) and Hansen (1985), according to which this parameter is not restricted by evidence on the Frisch elasticity.

2.1 Private economy

2.1.1 Households

We suppose there is a large number of identical households. The representative household solves the following problem:

$$\max_{\{C_t,H_t,B_{t+1}\}} E_0 \sum_{t=0}^{\infty} \beta^t \left(\log C_t - \frac{H_t^{1+\phi}}{1+\phi} \right), 0 < \beta < 1, \phi \geq 0, \tag{1}$$

subject to

$$P_t C_t + B_{t+1} \leq B_t R_{t-1} + W_t H_t + \text{Transfers and profits}_t. \tag{2}$$

Here, C_t and H_t denote household consumption and market work, respectively. In Eq. (2), B_{t+1} denotes the quantity of a nominal bond purchased by the household in period t and R_t denotes the one-period gross nominal rate of interest on a bond purchased in period t. Finally, W_t denotes the competitively determined nominal wage rate. The parameter, ϕ, is discussed in Section 2.3.

The representative household equates the marginal cost of working, in consumption units, with the marginal benefit, the real wage:

$$C_t H_t^{\phi} = \frac{W_t}{P_t}. \tag{3}$$

The representative household also equates the utility cost of the consumption foregone in acquiring a bond with the corresponding benefit:

$$\frac{1}{C_t} = \beta E_t \frac{1}{C_{t+1}} \frac{R_t}{\pi_{t+1}}. \tag{4}$$

Here, π_{t+1} denotes the gross rate of inflation from t to $t+1$.

2.1.2 Firms

A key feature of the New Keynesian model is its assumption that there are price-setting frictions. These frictions are introduced to accommodate the evidence of inertia in

aggregate inflation. Obviously, the presence of price-setting frictions requires that firms have the power to set prices, and this in turn requires the presence of monopoly power. A challenge is to create an environment in which there is monopoly power, without contradicting the obvious fact that actual economies have a very large number of firms. The Dixit-Stiglitz framework of production handles this challenge very nicely, because it has a very large number of price-setting monopolist firms. In particular, gross output is produced using a representative, competitive firm using the following technology:

$$Y_t = \left(\int_0^1 Y_{i,t}^{\frac{1}{\lambda_f}} di \right)^{\lambda_f}, \lambda_f > 1, \tag{5}$$

where λ_f governs the degree of substitution between the different inputs. The representative firm takes the price of gross output, P_t, and the price of intermediate inputs, $P_{i,t}$, as given. Profit maximization leads to the following first-order condition:

$$Y_{i,t} = Y_t \left(\frac{P_{i,t}}{P_t} \right)^{-\frac{\lambda_f}{\lambda_f - 1}}. \tag{6}$$

Substituting Eq. (6) into Eq. (5) yields the following relation between the aggregate price level and the prices of intermediate goods:

$$P_t = \left(\int_0^1 P_{i,t}^{-\frac{1}{\lambda_f - 1}} di \right)^{-(\lambda_f - 1)}. \tag{7}$$

The i^{th} intermediate good is produced by a single monopolist, who takes Eq. (6) as its demand curve. The value of λ_f determines how much monopoly power the i^{th} producer has. If λ_f is large, then intermediate goods are poor substitutes for each other, and the monopoly supplier of good i has a lot of market power. Consistent with this, note that if λ_f is large, then the demand for $Y_{i,t}$ is relatively price inelastic (see Eq. 6). If λ_f is close to unity, so that each $Y_{i,t}$ is almost a perfect substitute for $Y_{j,t}$, $j \neq i$, then the i^{th} firm faces a demand curve that is almost perfectly elastic. In this case, the firm has virtually no market power.

The production function of the i^{th} monopolist is:

$$Y_{i,t} = z_t H_{i,t}^{\gamma} I_{i,t}^{1-\gamma}, 0 < \gamma \leq 1, \tag{8}$$

where z_t is a technology shock whose stochastic properties are specified below. Here, $H_{i,t}$, denotes the level of employment by the i^{th} monopolist. We follow Basu (1995) in supposing that the i^{th} monopolist uses the quantity of materials, $I_{i,t}$, as inputs to production. The materials, $I_{i,t}$, are converted one-for-one from Y_t in Eq. (5). For $\gamma < 1$, each intermediate good producer in effect uses the output of all the other intermediate produces as input. When $\gamma = 1$, then materials inputs are not used in production.

The nominal marginal cost of the intermediate good producer is the following Cobb-Douglas function of the price of its two inputs:

$$\text{marginal cost}_t = \left(\frac{\bar{P}_t}{1-\gamma}\right)^{1-\gamma}\left(\frac{\bar{W}_t}{\gamma}\right)^{\gamma}\frac{1}{z_t}.$$

Here, $\bar{W}$ and $\bar{P}$ are the effective prices of $H_{i,t}$, and $I_{i,t}$, respectively:

$$\begin{aligned}\bar{W}_t &= (1-\nu_t)(1-\psi+\psi R_t)W_t\\ \bar{P}_t &= (1-\nu_t)(1-\psi+\psi R_t)P_t.\end{aligned} \tag{9}$$

In this expression, ν_t denotes a subsidy to intermediate good firms and the term involving the interest rate reflects the presence of a "working capital channel." For example, $\psi = 1$ corresponds to the case where the full amount of the cost of labor and materials must be financed at the beginning of the period. When $\psi = 0$, no advanced financing is required. A key variable in the model is the ratio of nominal marginal cost to the price of gross output, P_t:

$$s_t = (1-\nu_t)\left(\frac{1}{1-\gamma}\right)^{1-\gamma}\left(\frac{\bar{w}_t}{\gamma}\right)^{\gamma}(1-\psi+\psi R_t), \tag{10}$$

where $\bar{w}_t$ denotes the scaled real wage rate:

$$\bar{w}_t \equiv \frac{W_t}{z_t^{\frac{1}{\gamma}}Pt}. \tag{11}$$

If intermediate good firms faced no price-setting frictions, they would all set their price as a fixed markup over nominal marginal cost:

$$\lambda_f P_t s_t. \tag{12}$$

In fact, we assume there are price-setting frictions along the lines proposed by Calvo (1983). An intermediate firm can set its price optimally with probability $1-\xi_p$, and with probability ξ_p it must keep its price unchanged relative to what it was in the previous period:

$$P_{i,t} = P_{i,t-1}.$$

Consider the $1-\xi_p$ intermediate good firms that are able to set their prices optimally in period t. There are no state variables in the intermediate good firm problem and all the firms face the same demand curve. As a result, all firms able to optimize their prices in period t choose the same price, which we denote by $\widetilde{P}_t$. It is clear that optimizing firms do not set $\widetilde{P}_t$ equal to Eq. (12). Setting $\widetilde{P}_t$ to Eq. (12) would be optimal from the perspective of the current period, but it does not take into account the possibility that the firm may be stuck with $\widetilde{P}_t$ for several periods into the future. Instead, the intermediate good firms that have an opportunity to reoptimize their price in the current period, do so to solve:

$$\max_{\widetilde{P}_t} E_t\sum_{j=0}^{\infty}(\xi_p\beta)^j \upsilon_{t+j}\left(\widetilde{P}_t Y_{i,t+j} - P_{t+j}s_{t+j}Y_{i,t+j}\right), \tag{13}$$

subject to the demand curve, Eq. (6), and the definition of marginal cost, Eq. (10). In Eq. (13), $\beta^j \upsilon_{t+j}$ is the multiplier on the household's nominal period $t+j$ budget constraint. Because they are the owners of the intermediate good firms, households are the recipients of firm profits. In this way, it is natural that the firm should weigh profits in different dates and states of nature using $\beta^j \upsilon_{t+j}$. Intermediate good firms take υ_{t+j} as given. The nature of the family's preferences, Eq. (1), implies:

$$\upsilon_{t+j} = \frac{1}{P_{t+j} C_{t+j}}.$$

In Eq. (13) the presence of ξ_p reflects that intermediate good firms are only concerned with future scenarios in which they are not able to reoptimize the price chosen in period t.

The first-order condition associated with Eq. (13) is

$$\widetilde{p}_t = \frac{E_t \sum_{j=0}^{\infty} (\beta \xi_p)^j (X_{t,j})^{-\frac{\lambda_f}{\lambda_f - 1}} \lambda_f s_{t+j}}{E_t \sum_{j=0}^{\infty} (\beta \xi_p)^j (X_{i,j})^{-\frac{1}{\lambda_f - 1}}} = \frac{K_t^f}{F_t^f}, \tag{14}$$

where K_t^f and F_t^f denote the numerator and denominator of the ratio after the first equality, respectively. Also,

$$\widetilde{p}_t \equiv \frac{\widetilde{P}_t}{P_t}, X_{t,j} \equiv \begin{cases} \dfrac{1}{\pi_{t+j} \cdots \pi_{t+j}} & j > 0 \\ 1 & j = 0 \end{cases}.$$

Not surprisingly, Eq. (14) implies $\widetilde{P}_t$ is set to Eq. (12) when $\xi_p = 0$. When $\xi_p > 0$, optimizing firms set their prices so that Eq. (12) is satisfied on average. It is useful to write the numerator and denominator in Eq. (14) in recursive form. Thus,

$$K_t^f = \lambda_f s_t + \beta \xi_p E_t \pi_{t+1}^{\frac{\lambda_f}{\lambda_f - 1}} K_{t+1}^f, \tag{15}$$

$$F_t^f = 1 + \beta \xi_p E_t \pi_{t+1}^{\frac{1}{\lambda_f - 1}} F_{t+1}^f. \tag{16}$$

Expression (7) simplifies when we take into account that (i) the $1 - \xi_p$ intermediate good firms that set their price optimally all set it to $\widetilde{P}_t$ and (ii) the ξ_p firms that cannot reset their price are selected at random from the set of all firms. Doing so,

$$\widetilde{p}_t = \left[\frac{1 - \xi_p \pi_t^{\frac{1}{\lambda_f - 1}}}{1 - \xi_p} \right]^{-(\lambda_f - 1)}. \tag{17}$$

It is convenient to use Eq. (17) to eliminate $\widetilde{p}_t$ in Eq. (14):

$$K_t^f = F_t^f \left(\frac{1 - \xi_p \pi_t^{\frac{1}{\lambda_f - 1}}}{1 - \xi_p} \right)^{-(\lambda_f - 1)}. \tag{18}$$

When $\gamma < 1$, cost minimization by the i^{th} intermediate good producer leads it to equate the relative price of its labor and materials inputs to the corresponding relative marginal productivities:

$$\frac{\bar{W}_t}{\bar{P}_t} = \frac{W_t}{P_t} = \frac{\gamma}{1-\gamma} \frac{I_{i,t}}{H_{i,t}} = \frac{\gamma}{1-\gamma} \frac{I_t}{H_t}. \tag{19}$$

Evidently, each firm uses the same ratio of inputs, regardless of its output price, $P_{i,t}$.

2.1.3 Aggregate resources and the private sector equilibrium conditions

A notable feature of the New Keynesian model is the absence of an aggregate production function. That is, given information about aggregate inputs and technology, it is not possible to say what aggregate output, Y_t, is. This is because Y_t also depends on how inputs are distributed among the various intermediate good producers. For a given amount of aggregate inputs, Y_t is maximized by distributing the inputs equally across producers. An unequal distribution of inputs results in a lower level of Y_t. In the New Keynesian model with Calvo price frictions, resources are unequally allocated across intermediate good firms if, and only if, $P_{i,t}$ differs across i. Price dispersion in the model is caused by the interaction of inflation with price-setting frictions. With price dispersion, the price mechanism ceases to allocate resources efficiently, as too much production is done in firms with low prices and too little in the firms with high prices. Yun (1996) derived a very simple formula that characterizes the loss of output due to price dispersion. We re-derive the analog of Yun's (1996) formula that is relevant for our setting.

Let Y_t^* denote the unweighted integral of gross output across intermediate good producers:

$$Y_t^* \equiv \int_0^1 Y_{i,t} di = \int_0^1 z_t \left(\frac{H_{i,t}}{I_{i,t}} \right)^{\gamma} I_{i,t} di = z_t \left(\frac{H_t}{I_t} \right)^{\gamma} I_t = z_t H_t^{\gamma} I_t^{1-\gamma}.$$

Here, we have used linear homogeneity of the production function, as well as the result in Eq. (19), that all intermediate good producers use the same labor to materials ratio. An alternative representation of Y_t^* makes use of the demand curve, Eq. (6):

$$Y_t^* = Y_t \int_0^1 \left(\frac{P_{i,t}}{P_t} \right)^{-\frac{\lambda_f}{\lambda_f - 1}} di = Y_t P_t^{\frac{\lambda_f}{\lambda_f - 1}} \int_0^1 (P_{i,t})^{-\frac{\lambda_f}{\lambda_f - 1}} di = Y_t P_t^{\frac{\lambda_f}{\lambda_f - 1}} (P_t^*)^{-\frac{\lambda_f}{\lambda_f - 1}}. \tag{20}$$

Thus,

$$Y_t = p_t^* z_t H_t^{\gamma} I_t^{1-\gamma},$$

where

$$p_t^* \equiv \left(\frac{P_t^*}{P_t}\right)^{\frac{\lambda_f}{\lambda_f - 1}}. \tag{21}$$

Here, $P_t^* \leq 1$ denotes Yun's (1996) measure of the output lost due to price dispersion. From Eq. (20),

$$P_t^* = \left[\int_0^1 (P_{i,t})^{-\frac{\lambda_f}{\lambda_f - 1}} di\right]^{-\frac{\lambda_f - 1}{\lambda_f}}. \tag{22}$$

According to Eq. (21), P_t^* is a monotone function of the ratio of two different weighted averages of intermediate good prices. The ratio of these two weighted averages can only be at its maximum of unity if all prices are the same.[3]

Taking into account observations (i) and (ii) after Eq. (16), Eq. (22) reduces (after dividing by P_t and taking into account Eq. 21) to:

$$p_t^* = \left[(1 - \xi_p)\left(\frac{1 - \xi_p \pi_t^{\frac{1}{\lambda_f - 1}}}{1 - \xi_p}\right)^{\lambda_f} + \xi_p \frac{\pi_t^{\frac{\lambda_f}{\lambda_f - 1}}}{p_{t-1}^*}\right]^{-1}. \tag{23}$$

According to Eq. (23), there is price dispersion in the current period if there was dispersion in the previous period and/or if there is a current shock to dispersion. Such a shock must operate through the aggregate rate of inflation.

We conclude that the relation between aggregate inputs and gross output is given by:

$$C_t + I_t = p_t^* z_t H_t^{\gamma} I_t^{1-\gamma}. \tag{24}$$

Here, $C_t + I_t$ represents total gross output, while C_t represents value added.

The private sector equilibrium conditions of the model are Eqs. (3), (4), (10), (15), (16), (18), (19), (23), and (24). This represents 9 equations in the following 11 unknowns:

$$C_t, H_t, I_t, R_t, \pi_t, p_t^*, K_t^f, F_t^f, \frac{W_t}{P_t}, s_t, v_t. \tag{25}$$

[3] The distortion, p_t^*, is of interest in its own right. It is a sort of "endogenous Solow residual" of the kind called for by Prescott (1998). Whether the magnitude of fluctuations in p_t^* are quantitatively important given the actual price dispersion in data is something that deserves exploration. A difficulty that must be overcome, in such an exploration, is determining what the benchmark efficient dispersion of prices is in the data. In the model it is efficient for all prices to be exactly the same, but that is obviously only a convenient normalization.

As it stands, the system is underdetermined. This is not surprising, since we have said nothing about monetary policy or how v_t is determined. We turn to this in the following section.

2.2 Log-linearized equilibrium with Taylor rule

We log-linearize the equilibrium conditions of the model about its nonstochastic steady state. We assume that monetary policy is governed by a Taylor rule, which responds to the deviation between actual inflation and a zero inflation target. As a result, inflation is zero in the nonstochastic steady state. In addition, we suppose that the intermediate good subsidy, v_t, is set to the constant value that causes the price of goods to equal the social marginal cost of production in steady state. To see what this implies for v_t, recall that in steady state firms set their price as a markup, λ_f, over marginal cost. That is, they equate the object in Eq. (12) to P_t, so that:

$$\lambda_f s = 1.$$

Using Eq. (10) to substitute out for the steady state value of s, the latter expression reduces, in steady state, to:

$$\lambda_t(1-v)(1-\psi+\psi R)\left[\left(\frac{1}{1-\gamma}\right)^{1-\gamma}\left(\frac{\bar{w}}{\gamma}\right)^{\gamma}\right]=1.$$

Because we assume competitive labor markets, the object in square brackets is the ratio of social marginal cost to price. As a result, it is socially efficient for this expression to equal unity. This is accomplished in the steady state by setting v as follows:

$$1-v=\frac{1}{\lambda_f(1-\psi+\psi R)}. \tag{26}$$

Our treatment of policy implies that the steady-state allocations of our model economy are efficient in the sense that they coincide with the solution to a particular planning problem. To define this problem, it is convenient to adopt the following scaling of variables:

$$c_t \equiv \frac{C_t}{z_t^{1/\gamma}}, i_t \equiv \frac{I_t}{z_t^{1/\gamma}}. \tag{27}$$

The planning problem is:

$$\max_{\{c_t,H_t,i_t\}} E_0\sum_{t=0}^{\infty}\beta^t\left[\log c_t-\frac{H_t^{1+\phi}}{1+\phi}\right], \text{subject to } c_t+i_t=H_t^{\gamma}i_t^{1-\gamma}. \tag{28}$$

The problem, (28), is that of a planner who allocates resources efficiently across intermediate goods and who does not permit monopoly power distortions. Because there is no

state variable in the problem, it is obvious that the choice variables that solve Eq. (28) are constant over time. This implies that the C_t and I_t that solve the planning problem are a fixed proportion of $z_t^{1/\gamma}$ over time. It turns out that the allocations that solve Eq. (28) also solve the Ramsey optimal policy problem of maximizing Eq. (1) with respect to the 11 variables listed in Eq. (25) subject to the 9 equations listed before Eq. (25).[4]

Because inflation, π_t, fluctuates in equilibrium, Eq. (23) suggests that p_t^* fluctuates too. It turns out, however, that p_t^* is constant to a first-order approximation. To see this, note that the absence of inflation in the steady state also guarantees there is no price dispersion in steady state in the sense that p_t^* is at its maximal value of unity (see Eq. 23). With p_t^* at its maximum in steady state, small perturbations have a zero first-order impact on p_t^*. This can be seen by noting that π_t is absent from the log-linear expansion of Eq. (23) about $p_t^* = 1$:

$$\hat{p}_t^* = \xi_p \hat{p}_{t-1}^*. \tag{29}$$

Here, a hat over a variable indicates:

$$\hat{\varrho}_t = \frac{d\varrho_t}{\varrho},$$

where ϱ denotes the steady state of the variable, ϱ_t, and $d\varrho_t = \varrho_t - \varrho$ denotes a small perturbation in ϱ_t from steady state. We suppose that in the initial period, $\hat{p}_{t-1}^* = 0$, so that, to a first-order approximation, $\hat{p}_t^* = 0$ for all t.

Log-linearizing Eqs. (15), (16), and (18) we obtain the usual representation of the Phillips curve:

$$\hat{\pi}_t = \frac{(1-\beta\xi_p)(1-\xi_p)}{\xi_p}\hat{s}_t + \beta E_t \hat{\pi}_{t+1}. \tag{30}$$

Combining Eq. (3) with Eq. (10), taking into account Eq. (27) and the setting of v in Eq. (26), real marginal cost is:

$$s_t = \frac{1}{\lambda_f}\frac{1-\psi+\psi R_t}{1-\psi+\psi R}\left(\frac{1}{1-\gamma}\right)^{1-\gamma}\left(\frac{c_t H_t^{\phi}}{\gamma}\right)^{\gamma}.$$

Then,

$$\hat{s}_t = \gamma(\phi\hat{H}_t + \hat{c}_t) + \frac{\psi}{(1-\psi)\beta+\psi}\hat{R}_t. \tag{31}$$

Substituting out for the real wage in Eq. (19) using Eq. (3) and applying Eq. (27),

[4] The statement in the text is strictly true only in the case where the initial distortion in prices is zero, that is $p_{t-1}^* = 1$. If this condition does not hold, then the statement still holds asymptotically and may even hold as an approximation after a small number of periods.

$$H_t^{\phi+1} c_t = \frac{\gamma}{1-\gamma} i_t. \tag{32}$$

Similarly, scaling Eq.(24):

$$c_t + i_t = H_t^{\gamma} i_t^{1-\gamma}.$$

Using Eq. (32) to substitute out for i_t in the above expression, we obtain:

$$c_t + \frac{1-\gamma}{\gamma} H_t^{\phi+1} c_t = H_t^{\gamma} \left[\frac{1-\gamma}{\gamma} H_t^{\phi+1} c_t \right]^{1-\gamma}.$$

Log-linearizing this expression around the steady state implies, after some algebra,

$$\hat{c}_t = \hat{H}_t. \tag{33}$$

Substituting the latter into Eq. (31), we obtain:

$$\hat{s}_t = \gamma(1+\phi)\hat{c}_t + \frac{\psi}{(1-\psi)\beta + \psi} \hat{R}_t. \tag{34}$$

In Eq. (34), $\hat{c}_t$ is the percent deviation of c_t from its steady-state value. Since this steady-state value coincides with the constant c_t that solves Eq. (28) for each t, $\hat{c}_t$ also corresponds to the output gap. The notation we use to denote the output gap is x_t. Using this notation for the output gap and substituting out $\hat{s}_t$ in the Phillips curve, we obtain:

$$\hat{\pi}_t = \kappa_p \left[\gamma(1+\phi) x_t + \frac{\psi}{(1-\psi)\beta + \psi} \hat{R}_t \right] + \beta E_t \hat{\pi}_{t+1}, \tag{35}$$

where

$$\kappa_p \equiv \frac{(1-\beta\xi_p)(1-\xi_p)}{\xi_p}.$$

When $\gamma = 1$ and $\psi = 0$, Eq. (35) reduces to the "Phillips curve" in the classic New Keynesian model. When materials are an important factor of production, so that γ is small, then a given jump in the output gap, x_t, has a smaller impact on inflation. The reason is that in this case the aggregate price index is part of the input cost for intermediate good producers. So, a small price response to a given output gap is an equilibrium because individual intermediate good firms have less of an incentive to raise their prices in this case. With $\psi > 0$, Eq. (35) indicates that a jump in the interest rate drives up prices. This is because with an active working capital channel a rise in the interest rate drives up marginal cost.[5]

[5] Equation (35) resembles equation (13) in Ravenna and Walsh (2006), except that we also allow for materials inputs, i.e., $\gamma < 1$.

Now consider the intertemporal Euler equation. Expressing (4) in terms of scaled variables,

$$1 = E_t \frac{\beta c_t}{c_{t+1}\mu_{z,t+1}^{\frac{1}{\gamma}}} \frac{R_t}{\pi_{t+1}}, \mu_{z,t+1} \equiv \frac{z_{t+1}}{z_t}.$$

Log-linearly expanding about steady state and recalling that $\hat{c}_t$ corresponds to the output gap:

$$0 = E_t\left[x_t - x_{t+1} - \frac{1}{\gamma}\hat{\mu}_{z,t+1} + \hat{R}_t - \hat{\pi}_{t+1}\right],$$

or,

$$x_t = E_t\left[x_{t+1} - \left(\hat{R}_t - \hat{\pi}_{t+1} - \hat{R}_t^*\right)\right], \tag{36}$$

Where

$$\hat{R}_t^* \equiv \frac{1}{\gamma} E_t \hat{\mu}_{z,t+1}. \tag{37}$$

We suppose that monetary policy, when linearized about steady state, is characterized by the following Taylor rule:

$$\hat{R}_t = r_\pi E_t \hat{\pi}_{t+1} + r_x x_t. \tag{38}$$

The equilibrium of the log-linearly expanded economy is given by Eq. (35) to (38).

2.3 Frisch labor supply elasticity

The magnitude of the parameter, ϕ, in the household utility function plays an important role in the analysis in later sections. This parameter has been the focus of much debate in macroeconomics. Note from Eq. (3) that the elasticity of H_t with respect to the real wage, holding C_t constant, is $1/\phi$. The condition, "holding C_t constant," could mean that the elasticity refers to the response of H_t to a change in the real wage that is of very short duration, so short that the household's wealth — and, hence, consumption — is left unaffected. Alternatively, the elasticity could refer to the response of H_t to a change in the real wage that is associated with an offsetting lump-sum transfer payment that keeps wealth unchanged. The debate about ϕ centers on the interpretation of H_t. Under one interpretation, H_t represents the amount of hours worked by a typical person in the labor force. With this interpretation, $1/\phi$ is the Frisch labor supply elasticity.[6] This is perhaps the most straightforward interpretation of $1/\phi$ given our

[6] The Frisch labor supply elasticity refers to the substitution effect associated with a change in the wage rate. It is the percent change in a person's labor supply in response to a change in the real wage, holding the marginal utility of consumption fixed. Throughout this chapter, we assume that utility is additively separable in consumption and leisure, so that constancy of the marginal utility of consumption translates into constancy of consumption.

assumption that the economy is populated by identical households, in which H_t is the labor effort of the typical household. An alternative interpretation of H_t is that it represents the number of people working, and that $1/\phi$ measures the elasticity with which marginal people substitute in and out of employment in response to a change in the wage. Under this interpretation, $1/\phi$ need not correspond to the labor supply elasticity of any particular person. The two different interpretations of H_t give rise to very different views about how data ought to be used to restrict the value of ϕ.

There is an influential labor market literature that estimates the Frisch labor supply elasticity using household level data. The general finding is that, although the Frisch elasticity varies somewhat across different types of people, on the whole the elasticities are very small. Some have interpreted this to mean that only large values of ϕ (say, larger than unity) are consistent with the data. Initially, this interpretation was widely accepted by macroeconomists. However, the interpretation gave rise to a puzzle for equilibrium models of the business cycle. Over the business cycle, employment fluctuates a great deal more than real wages. When viewed through the prism of equilibrium models the aggregate data appeared to suggest that people respond elastically to changes in the wage. But, this seemed inconsistent with the microeconomic evidence that individual labor supply elasticities are in fact small. At the present time, a consensus is emerging that what initially appeared to be a conflict between micro and macro data is really no conflict at all. The idea is that the Frisch elasticity in the micro data and the labor supply elasticity in the macro data represent at best distantly related objects.

It is well known that much of the business cycle variation in employment reflects changes in the quantity of people working, not in the number of hours worked by a typical household. Beginning at least with the work of Rogerson (1988) and Hansen (1985), it has been argued that even if the individual's labor supply elasticity is *zero* over most values of the wage, aggregate employment could nevertheless respond highly elastically to small changes in the real wage. This can occur if there are many people who are near the margin between working in the market and devoting their time to other activities. An example is a spouse who is doing productive work in the home, and yet who might be tempted by a small rise in the market wage to substitute into the market. Another example is a teenager who is close to the margin between pursuing additional education and working, who could be induced to switch to working by a small rise in the wage. Finally, there is the elderly person who might be induced by a small rise in the market wage to delay retirement. These examples suggest that aggregate employment might fluctuate substantially in response to small changes in the real wage, even if the individual household's Frisch elasticity of labor supply is zero over all values of the wage, except the one value that induces a shift in or out of the labor market.[7]

[7] See Rogerson and Wallenius (2009) for additional discussion and analysis.

The ideas in the previous paragraphs can be illustrated in our model. We adopt the technically convenient assumption that the household has a large number of members, one for each of the points on the line bounded by 0 and 1.[8] In addition, we assume that a household member only has the option to work full time or not at all. A household member's Frisch labor supply elasticity is zero for almost all values of the wage. Let $l \in [0, 1]$ index a particular member in the family. Suppose this member enjoys the following utility if employed:

$$\log C_t - l^{\phi}, \phi > 0,$$

and the following utility if not employed:

$$\log C_t.$$

Household members are ordered according to their degree of aversion to work. Those with high values of l have a high aversion (e.g., small children, and elderly or chronically ill people) to work, and those with l near zero have very little aversion. We suppose that household decisions are made on a utilitarian basis, in a way that maximizes the equally weighted integral of utility across all household members. Under these circumstances, efficiency dictates that all members receive the same level of consumption, whether employed or not. In addition, if H_t members are to be employed, then those with $0 \leq l \leq H_t$ should work and those with $l > H_t$ should not. For a household with consumption, C_t, and employment, H_t, utility is, after integrating over all $l \in [0, 1]$:

$$\log C_t - \frac{H_t^{1+\phi}}{1+\phi}, \tag{39}$$

which coincides with the period utility function in Eq. (1). Under this interpretation of the utility function, Eq. (3) remains the relevant first-order condition for labor. In this case, given the wage, W_t/P_t, the household sends out a number of members, H_t, to work until the utility cost of work for the marginal worker, H_t^{ϕ}, is equated to the corresponding utility benefit to the household, $(W_t/P_t)/C_t$.

Note that under this interpretation of the utility function, H_t denotes a quantity of workers and ϕ dictates the elasticity with which different members of the households enter or leave employment in response to shocks. The case in which ϕ is large corresponds to the case where household members differ relatively sharply in terms of their aversion to work. In this case there are not many members with disutility of work close to that of the marginal worker. As a result, a given change in the wage induces only a small change in employment. If ϕ is very small, then there is a large number of

[8] Our approach is most similar to the approach of Gali (2010a), although it also resembles the approach taken in the recent work of Mulligan (2001) and Krusell, Mukoyama, Rogerson, and Sahin (2008).

household members close to indifferent between working and not working, and so a small change in the real wage elicits a large labor supply response.

Given that most of the business cycle variation in the labor input is in the form of numbers of people employed, we think the most sensible interpretation of H_t is that it measures numbers of people working. Accordingly, $1/\phi$ is not to be interpreted as a Frisch elasticity, which we instead assume to be zero.

3. SIMPLE MODEL: SOME IMPLICATIONS FOR MONETARY POLICY

Monetary DSGE models have been used to gain insight into a variety of issues that are important for monetary policy. We discuss some of these issues using variants of the simple model developed in the previous section. A key feature of that model is that it is flexible, and can be adjusted to suit different questions and points of view. The classic New Keynesian model, the one with no working capital channel and no materials inputs (i.e., $\gamma = 1$, $\psi = 0$) can be used to articulate the rationale for the Taylor principle. But variants of the New Keynesian framework can also be used to articulate challenges to that principle. Sections 3.1 and 3.2 below describe two such challenges. The fact that the New Keynesian framework can accommodate a variety of perspectives on important policy questions is an important strength. This is because the framework helps to clarify debates and to motivate econometric analyses so that data can be used to resolve those debates.[9]

Sections 3.3 and 3.4 below address the problem of estimating the output gap. The output gap is an important variable for policy analysis because it is a measure of the efficiency with which economic resources are allocated. In addition, New Keynesian models imply that the output gap is an important determinant of inflation, a variable of particular concern to monetary policymakers. We define the output gap as the percent deviation between actual output and potential output, where potential output is output in the Ramsey-efficient equilibrium.[10]

We use the classic New Keynesian model to study three ways of estimating the output gap. The first uses the structure of the simple New Keynesian model to estimate the output gap as a latent variable. The second approach modifies the New Keynesian model to include unemployment along the lines indicated by CTW. This modification of the model allows us to investigate the information content of the unemployment rate for the output gap. In addition, by showing one way that unemployment can be integrated into the model, the discussion represents another illustration of the versatility

[9] For example, the Chowdhury, Hoffmann, and Schabert (2006) and Ravenna and Walsh (2006) papers cited in the previous section, show how the assumptions of the New Keynesian model can be used to develop an empirical characterization of the importance of the working capital channel.

[10] In our model, the Ramsey-equilibrium turns out to be what is often called the "first-best equilibrium," the one that is not distorted by monopoly power or flexible prices.

of the New Keynesian framework.[11] The third approach which is studied in section 3.4 explores the HP filter as a device for estimating the output gap. In the course of the analysis, we illustrate the Bayesian limited information moment matching procedure discussed in the introduction.

3.1 Taylor principle

A key objective of monetary policy is the maintenance of low and stable inflation. The classic New Keynesian model defined by $\gamma = 1$ and $\psi = 0$ can be used to articulate the risk that inflation expectations might become self-fulfilling unless the monetary authorities adopt the appropriate monetary policy. The classic model can also be used to explain the widespread consensus that "appropriate monetary" policy means a monetary policy that embeds the Taylor Principle: a 1% rise in inflation should be met by a greater than 1% rise in the nominal interest rate. This subsection explains how the classic New Keynesian model rationalizes the wisdom of implementing the Taylor principle. However, when we incorporate the assumption of a working capital channel — particularly when the share of materials in gross output is as high as it is in the data — the Taylor principle becomes a *source* of instability. This is perhaps not surprising. When the working capital channel is strong, if the monetary authority raises the interest rate in response to rising inflation expectations, the resulting rise in costs produces the higher inflation that people expect.[12]

It is convenient to summarize the linearized equations of our model here:

$$\hat{R}_t^* = E_t \frac{1}{\gamma} \hat{\mu}_{z,t+1} \tag{40}$$

$$\hat{\pi}_t = \kappa_p \left[\gamma(1+\phi) x_t + \alpha_\psi \hat{R}_t \right] + \beta E_t \hat{\pi}_{t+1} \tag{41}$$

[11] For an alternative recent approach to the introduction of unemployment into a DSGE model, see Gali (2010a). Gali demonstrated that with a modest reinterpretation of variables, the standard DSGE model with sticky wages summarized in the next section contains a theory of unemployment. In the model of the labor market used there (it was proposed by Erceg et al. 2000) wages are set by a monopoly union. As a result, the wage rate is higher than the marginal cost of working. Under these circumstances, one can define the unemployed as the difference between the number of people actually working and the number of people that would be working if the cost of work for the marginal person were equated to the wage rate. Gali (2010b) showed how unemployment data can be used to help estimate the output gap, as we do here. The CTW and Gali models of unemployment are quite different. For example, in the text we analyze a version of the CTW model in which labor markets are perfectly competitive, so Gali's "monopoly power" concept of unemployment is zero in this model. In addition, the efficient level of unemployment in the sense that we use the term here, is zero in Gali's definition, but positive in our definition. This is because in our model, unemployment is an inevitable by-product of an activity that must be undertaken to find a job. For an extensive discussion of the differences between our model and Gali's, see Section F in the technical appendix to CTW, which can be found at http://faculty.wcas.northwestern.edu/~lchrist/research/Riksbank/technicalappendix.pdf.

[12] Bruckner and Schabert (2003) made an argument similar to ours, although they do not consider the impact of materials inputs, which we find to be important.

$$x_t = E_t\left[x_{t+1} - \left(\hat{R}_t - \hat{\pi}_{t+1} - \hat{R}_t^*\right)\right] \tag{42}$$

$$\hat{R}_t = r_\pi E_t \hat{\pi}_{t+1} + r_x x_t, \tag{43}$$

where

$$\alpha_\psi = \frac{\psi}{(1-\psi)\beta + \psi}.$$

The specification of the model is complete when we take a stand on the law of motion for the exogenous shock. We do this in the following subsections as needed.

We begin by reviewing the case for the Taylor principle using the classic New Keynesian model, with $\gamma = 1$, $\psi = 0$. We get to the heart of the argument using the deterministic version of the model, in which $\hat{R}_t^* \equiv 0$. In addition, it is convenient to suppose that monetary policy is characterized by $r_x = 0$. Throughout, we adopt the presumption that the only valid equilibria are paths for $\hat{\pi}_t$, $\hat{R}_t$ and x_t that converge to steady state; that is, 0.[13] Under these circumstances, Eqs. (41) and (42) can be solved forward as follows:

$$\hat{\pi}_t = \kappa_p\gamma(1+\phi)x_t + \beta\kappa_p\gamma(1+\phi)x_{t+1} + \beta^2\kappa_p\gamma(1+\phi)x_{t+2} + \ldots \tag{44}$$

and

$$x_t = -\left(\hat{R}_t - \hat{\pi}_{t+1}\right) - \left(\hat{R}_{t+1} - \hat{\pi}_{t+2}\right) - \left(\hat{R}_{t+2} - \hat{\pi}_{t+3}\right) - \ldots \tag{45}$$

In Eq. (45) we have used the fact that in our setting a path converges to zero if, and only if, it converges fast enough so that a sum like the one in Eq. (45) is well defined.[14] Equation (44) shows that inflation is a function of the present and future output gap. Equation (45) shows that the current output gap is a function of the long term real interest rate (i.e., the sum on the right of Eq. 45) in the model.

Under the Taylor principle, the classic New Keynesian model implies that a rise in inflation expectations launches a sequence of events that ultimately leads to a

[13] Although our presumption is standard, justifying it is harder than one might have thought. For example, Benhabib, Schmitt-Grohe, and Uribe (2002) presented examples in which some explosive paths for the linearized equilibrium conditions are symptomatic of perfectly sensible equilibria for the actual economy underlying the linear approximations. In these cases, focusing on the nonexplosive paths of the linearized economy may be valid after all if we imagine that monetary policy is a Taylor rule with a particular escape clause. The escape clause specifies that in the event the economy threatens to follow an explosive path, the monetary authority commits to switch to a monetary policy of targeting the money growth rate. There are examples of monetary models in which the escape clause monetary policy justifies the type of equilibrium selection we adopt in the text (see Benhabib et al. 2002 and Christiano & Rostagno, 2001 for further discussion). For a more recent debate about the validity of the equilibrium selection adopted in the text, see McCallum (2009) and Cochrane (2009) and the references they cite.

[14] The reason for this can be seen below, where we show that the solution to this equation is a linear combination of terms like $a\lambda^t$. Such an expression converges to zero if, and only if, it is also summable.

moderation in actual inflation. Seeing this moderation in actual inflation, people's higher inflation expectations would quickly dissipate before they could be a source of economic instability. The way this works is that the rise in the real rate of interest slows spending, causing the output gap to shrink (see Eq. 45). The fall in actual inflation occurs as the reduction in output reduces pressure on resources and drives down the marginal cost of production (see Eq. 41). Strictly speaking, we have just described a rationale for the Taylor principle that is based on learning (for a formal discussion, see McCallum, 2009). Under rational expectations, the posited rise in inflation expectations would not occur in the first place if policy obeys the Taylor principle.

A similar argument shows that if the monetary authority does not obey the Taylor principle, that is, $r_\pi < 1$, then a rise in inflation expectations can be self-fulfilling. This is not surprising, since in this case the rise in expected inflation is associated with a fall in the real interest rate. According to Eq. (45) this produces a rise in the output gap. By raising marginal cost, the Phillips curve, (Eq. 44), implies that actual inflation rises. Seeing higher actual inflation, people's higher inflation expectations are confirmed. In this way, with $r_\pi < 1$ a rise in inflation expectations becomes self-fulfilling by triggering a boom in output and actual inflation. It is easy to see that with $r_\pi < 1$ many equilibria are possible. A drop in inflation expectations can cause a fall in output and inflation. Inflation expectations could be random, causing random fluctuations between booms and recessions.[15]

In this way, the classic New Keynesian model has been used to articulate the idea that the Taylor principle promotes stability, while absence of the Taylor principle makes the economy vulnerable to fluctuations in self-fulfilling expectations.

The preceding results are particularly easy to establish formally under the assumption of rational expectations. We continue to maintain the simplifying assumption, $r_x = 0$. We reduce the model to a single second order difference equation in inflation. Substitute out for $\hat{R}_t$ in Eqs. (41) and (42) using Eq. (43). Then, solve Eq. (41) for x_t and use this to substitute out for x_t in Eq. (42). These operations result in the following second-order difference equation in $\hat{\pi}_t$:

$$\hat{\pi}_t + [\kappa_p\gamma(1+\phi)(r_\pi - 1) - (\kappa_p\alpha_\psi r_\pi + \beta) - 1]\hat{\pi}_{t+1} + (\kappa_p\alpha_\psi r_\pi + \beta)\hat{\pi}_{t+2} = 0.$$

[15] Clarida, Gali, and Gertler (1999; CGG) argued that the high inflation of the 1970s in many countries can be explained as reflecting the failure to respect the Taylor principle in the early 1970s. Christiano and Gust (2000) criticized this argument on the grounds that one did not observe a boom in employment in the 1970s. Christiano and Gust argued that even if one thought of the 1970s as also a time of bad technology shocks (fuel costs and commodity prices soared then), the CGG analysis predicts that employment should have boomed. Christiano and Gust presented an alternative model, a "limited participation" model, which has the same implications for the Taylor principle that the CGG model has. However, the Christiano and Gust model has a very different implication for what happens to real allocations in a self-fulfilling inflation episode. Because of the presence of an important working capital channel, the self-fulfilling inflation episode is associated with a recession in output and employment. Thus, Christiano and Gust concluded that the 1970s might well reflect the failure to implement the Taylor principle, but only if the analysis is done in a model different from the CGG model.

The general set of solutions to this difference equation can be written as follows:

$$\hat{\pi}_t = a_0\lambda_1^t + a_1\lambda_2^t,$$

for arbitrary a_0, a_1. Here, λ_i, $i = 1, 2$, are the roots of the following second-order polynomial:

$$1 + [\kappa_p\gamma(1+\phi)(r_\pi - 1) - (\kappa_p\alpha_\psi r_\pi + \beta + 1)]\lambda + (\kappa_p\alpha_\psi r_\pi + \beta)\lambda^2 = 0.$$

Thus, there is a two-dimensional space of solutions to the equilibrium conditions (i.e., one for each possible value of a_0 and a_1). We continue to apply our presumption that among these solutions, only the ones in which the variables converge to zero (i.e., to steady state) correspond to equilibria. Thus, uniqueness of equilibrium requires that both λ_1 and λ_2 be larger than unity in absolute value. In this case, the unique equilibrium is the solution associated with $a_0 = a_1 = 0$. If one or both of λ_i, $i = 1, 2$ are less than unity in absolute value, then there are many solutions to the equilibrium conditions that are equilibria. We can think of these equilibria as corresponding to different, self-fulfilling, expectations.

The following result can be established for the classic New Keynesian model, with $\gamma = 1$ and $\psi = 0$. The model economy has a unique equilibrium if, and only if $r_\pi > 1$ (see, e.g., Bullard & Mitra, 2002). This is consistent with the intuition about the Taylor principle discussed above.

We now reexamine the case for the Taylor principle when there is a working capital channel. The reason the Taylor principle works in the classic New Keynesian model is that a rise in the interest rate leads to a fall in inflation by curtailing aggregate spending. But, with a working capital channel, $\psi > 0$, an increase in the interest rate has a second effect. By raising marginal cost (see Eq. 41), a rise in the interest rate places upward pressure on inflation. If the working capital channel is strong enough, then monetary policy with $r_\pi > 1$ may "add fuel to the fire" when inflation expectations rise. The sharp rise in the nominal rate of interest in response to a rise in inflation expectations may actually cause the inflation that people expected. In this way the Taylor principle could actually be destabilizing. Of course, for this to be true requires that the working capital channel be strong enough. For a small enough working capital channel (i.e., small ψ) implementing the Taylor principle would still have the effect of inoculating the economy from destabilizing fluctuations in inflation expectations.

Whether the presence of the working capital channel overturns the wisdom of implementing the Taylor principle is a numerical question. We must assign values to the model parameters and investigate whether one or both of λ_1 and λ_2 are less than unity in absolute value. If this is the case, then implementing the Taylor principle does not stabilize inflation expectations. Throughout, we set:

$$\beta = 0.99, \xi_p = 0.75, r_\pi = 1.5.$$

The discount rate is 4%, at an annual rate and the value of ξ_p implies an average time between price reoptimization of one year. In addition, monetary policy is characterized by a strong commitment to the Taylor principle. We consider two values for the interest rate response to the output gap, $r_x = 0$ and $r_x = 0.1$. For robustness, we also consider a version of Eq. (43) in which the monetary authority reacts to current inflation.

We do not have a strong prior about the parameter, ϕ, that controls the disutility of labor (see Section 2.3), so we consider two values, $\phi = 1$ and $\phi = 0.1$. To have a sense of the appropriate value of γ, we follow Basu (1995). He argued, using manufacturing data, that the share of materials in gross output is roughly 1/2. Recall that the steady state of our model coincides with the solution to Eq. (28), so that

$$\frac{i}{c+i} = 1 - \gamma.$$

Thus, Basu's empirical finding suggests a value for γ close to 1/2.[16] The instrumental variables results in Ravenna and Walsh (2006) suggest that a value of the working capital share, ψ, in a neighborhood of unity is consistent with the data.

Figure 1 displays our results. The upper row of figures provides results for the case in Eq. (43), in which the policy authority reacts to the one-quarter-ahead expectation of inflation, $E_t\hat{\pi}_{t+1}$. The lower row of figures corresponds to the case where the policymaker responds instead to current inflation, $\hat{\pi}_t$. The horizontal and vertical axes indicate a range of values for γ and ψ, respectively. The gray areas correspond to the parameter values where one or both of λ_i, $i = 1, 2$ are less than unity in absolute value. Technically, the steady-state equilibrium of the economy is said to be "indeterminate" for parameterizations in the gray area. Intuitively, the gray area corresponds to parameterizations of our economy in which the Taylor principle does not stabilize inflation expectations. The white areas in the figures correspond to parameterizations where implementing the Taylor principle successfully stabilizes the economy.

Consider the 1,1 and 1,2 graphs in Figure 1 first. Note that in each case, $\psi = 0$ and $\gamma = 1$ are points in the white area, consistent with the earlier discussion. However, a very small increase in the value of ψ puts the model into the gray area. Moreover, this is true regardless of the value of γ. For these parameterizations the aggressive response of the interest rate to higher inflation expectations only produces the higher inflation that people anticipate. We can see in the 1,3 and 1,4 graphs of the first row, that $r_x > 0$ greatly reduces the extent of the gray area. Still, for $\gamma = 0.5$ and ψ near unity the model is in the gray area and implementing the Taylor principle would be counterproductive.

[16] Actually, this is a conservative estimate of γ . Had we not selected v to extinguish monopoly power in the steady state, our estimate of γ would have been lower. See Basu (1995) for more discussion of this point.

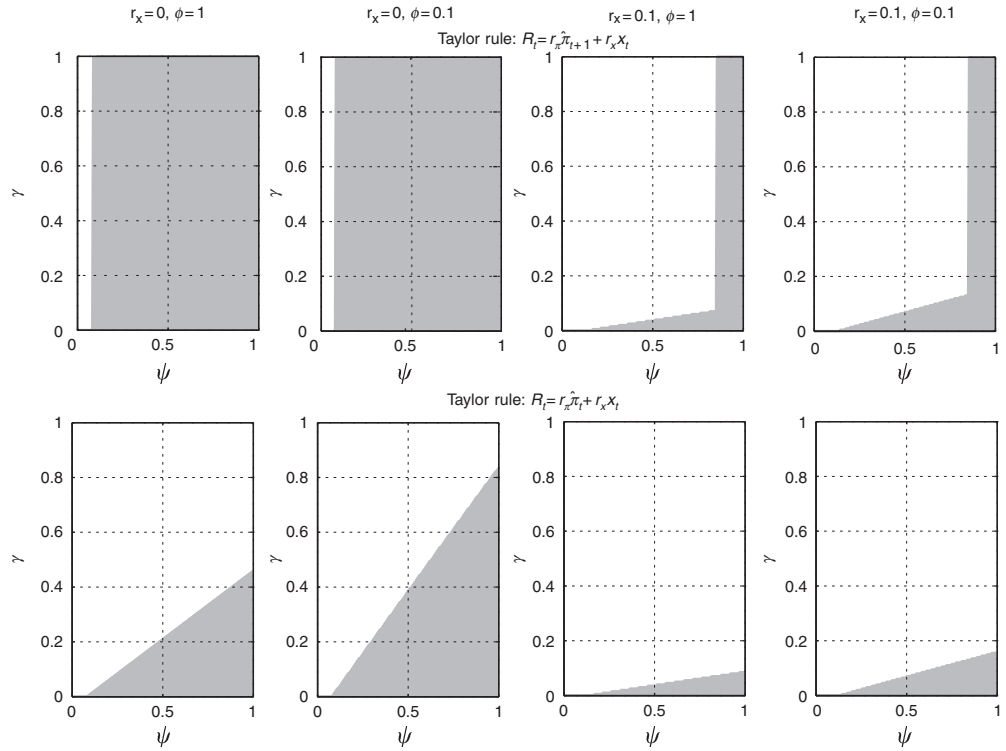

Figure 1 Indeterminacy regions for model with working capital channel and materials inputs. Note: Gray area is region of indeterminacy and white area is region of determinacy.

Now consider the bottom row of graphs. Note that if $\gamma = 1$ then the model is always in the determinacy region. That is, for the economy to be vulnerable to self-fulfilling expectations, it must not only be that there is a substantial working capital channel, but it must also be that materials are a substantial fraction of gross output. The 2,2 graph shows that with $\gamma = 0.5$, $\phi = 0.1$ and ψ above roughly 0.6, the model is in the gray area. When ϕ is substantially higher, the first graph from the left indicates that the gray area is smaller. Note that with $r_x > 0$, the gray area has almost shrunk to zero, according to the two last graphs.

We conclude from this analysis that in the presence of a working capital channel, sharply raising the interest rate in response to higher inflation could actually be counter-productive. This is more likely to be the case when the share of materials inputs in gross output is high. When this is so, one cannot rely exclusively on the Taylor principle to ensure stable inflation and output performance. In the example, responding strongly to the output gap (or, actual rather than expected inflation) could restore stability. However,

in practice the output gap is hard to measure.[17] At best, the policy authority can respond to variables that are correlated with the output gap. Studying the implications for determinacy of responding to such variables would be an interesting project, but would take us beyond the scope of this chapter. Still, the discussion illustrates how DSGE models can be useful for thinking about important monetary policy questions.

3.2 Monetary policy and inefficient booms

In recent years, there has been extensive discussion about the interaction of monetary policy and economic volatility, in particular, asset price volatility. Prior to the recent financial turmoil, a consensus had developed that monetary policy should not actively seek to stabilize asset prices. The view was that in any case, a serious commitment to inflation targeting — one that implements the Taylor principle — would stabilize asset markets automatically.[18] The idea is that an asset price boom is basically a demand boom, the presumption being that the boom is driven by optimism about the future, and not primarily by current actual developments. A boom that is driven by demand should, according to the conventional wisdom, raise production costs and, hence, inflation. The monetary authority that reacts vigorously to inflation then automatically raises interest rates and helps to stabilize asset prices.

When this scenario is evaluated in the classic New Keynesian model, we find that the boom is not necessarily associated with a rise in prices. In fact, if the optimism about the future concerns the expectations about cost saving new technologies, forward-looking price-setters may actually reduce their prices. This is the finding of Christiano, Ilut, Motto, and Rostagno (2008, 2010), which we briefly summarize here.

To capture the notion of optimism about the future, suppose that the time series representation of the log-level of technology is as follows:

$$\log z_t = \rho_z \log z_{t-1} + u_t, u_t = \varepsilon_t + \xi_{t-1}, \tag{46}$$

so that the steady state of z_t is unity. In Eq. (46), u_t is an iid shock, uncorrelated with past log z_t. The innovation in technology growth, u_t, is the sum of two orthogonal processes, ε_t and ξ_{t-1}. The time subscript on these two variables represents the date when they are known to private agents. Thus, at time $t-1$ agents become aware of a component of u_t, namely ξ_{t-1}. At time t they learn the rest, ε_t. For example, the initial "news" about u_t, ξ_{t-1}, could in principle be entirely false, as would be the case when $\varepsilon_t = -\xi_{t-1}$.

Substituting Eq. (46) into Eq. (40):

$$\hat{R}_t^* = E_t[\log z_{t+1} - \log z_t] = (\rho_z - 1)\log z_t + \xi_t, \tag{47}$$

[17] For further discussion of this point, see Sections 3.3 and 3.4.
[18] See Bernanke and Gertler (2000).

where $\gamma = 1$ since we now consider the classic New Keynesian model.[19] Our system of equilibrium conditions is Eq. (47) with Eqs. (41), (42) and (43). We set $\psi = 0$ (i.e., no working capital channel) and $r_x = 0$. We adopt the following parameter values:

$$\beta = 0.99, \phi = 1, r_x = 0, r_\pi = 1.5, \rho_z = 0.9, \xi_p = 0.75.$$

We perform a simulation in which news arrives in period t that technology will jump one percent in period $t + 1$, i.e., $\xi_t = 0.01$. The value of ε_t is set to zero. We find that hours worked in period t increases by 1%. This rise is entirely inefficient because in the first-best equilibrium hours it does not respond at all to a technology shock, whether it occurs in the present or it is expected to occur in the future (see Eq. 28). Interestingly, inflation falls in period t by 10 basis points, at an annual rate.[20] Current marginal cost does rise (see Eq. 34), but current inflation nevertheless falls because of the fall in expected future marginal costs.

The efficient monetary policy sets $\hat{R}_t = \hat{R}_t^*$ which, according to Eq. (47), means the interest rate should rise when a positive signal about the economy occurs. A policy that applies the Taylor principle in this example moves policy in exactly the wrong direction in response to ξ_t. By responding to the fall in inflation, policy not only does not raise the interest rate — as it should — but it actually reduces the interest rate in response to the fall in inflation. By reducing the interest rate in the period of a positive signal about the future, policy overstimulates the economy creating excessive volatility.

So, the classic New Keynesian model can be used to challenge the conventional wisdom that an inflation-fighting central bank automatically moderates economic volatility. But, is this just an abstract example without any relevance? In fact, the typical boom-bust episode is characterized by low or falling inflation (see Adalid & Detken, 2007). For example, during the U.S. booms of the 1920s and the 1980s and 1990s, inflation was low. This fact turns the conventional wisdom on its head and leads to a conclusion that matches our numerical example: an inflation-fighting central bank amplifies boom-bust episodes.

A full evaluation of the ideas in this subsection requires a more elaborate model, preferably one with financial variables such as the stock market. In this way, one could assess the impact on a broader set of variables in boom-bust episodes. In addition, one could evaluate what other variables the monetary authority might look at to avoid contributing to the type of volatility described in this example. We presume that it is not helpful to simply say that the monetary authority should set $\hat{R}_t = \hat{R}_t^*$, because in practice this may require more information than is actually available. A more fruitful

[19] To see why we replaced $\hat{\mu}_{z,t+1}$ in Eq. (40) by $\log z_{t+1} - \log z_t$, note first $\hat{\mu}_{z,t} = \frac{\mu_{z,t} - \mu_z}{\mu_z} = \mu_{z,t} - 1$, because in steady state $\mu_z \equiv z_t / z_{t-1} = 1/1 = 1$. Then, $1 + \hat{\mu}_{z,t} = \mu_{z,t}$. Take the log of both sides and note, $\log \mu_{z,t} = \log (1 + \hat{\mu}_{z,t}) \simeq \hat{\mu}_{z,t}$. But, $\log \mu_{z,t} = \log z_t - \log z_{t-1}$.

[20] Because inflation is zero in steady state, $\hat{\pi}_t = \pi_t - 1$. This was converted to annualized basis points by multiplying by 40,000.

approach may be to find variables that are correlated with $\hat{R}_t^*$, so that these may be included in the monetary policy rule. For further discussion of these issues, see Christiano et al. (2008).

3.3 Using unemployment to estimate the output gap

Here, we investigate the use of DSGE models to estimate the output gap as a latent variable. We explore the usefulness of including data on the rate of unemployment in this exercise. Section 3.3.1 describes a scalar statistic for characterizing the information content of the unemployment rate for the output gap, and Section 3.3.2 describes the model used in the analysis. As in the previous subsection, we work with a version of the classic New Keynesian model. In particular, we assume intermediate good producers do not use materials inputs or working capital.[21] We introduce unemployment into the model following the approach in CTW.

Section 3.3.3 describes how we use data to assign values to the model parameters. This section may be of independent interest because it shows how a moment-matching procedure like the one proposed in Christiano and Eichenbaum (1992a) can be recast in Bayesian terms. Section 3.3.4 presents our substantive results. Based on our simple estimated model with unemployment, we find that including unemployment has a substantial impact on our estimate of the output gap for the U.S. economy. We summarize our findings at the end of Section 3.3.4 where we also indicate several caveats to the analysis.

3.3.1 A measure of the information content of unemployment

As a benchmark, we compute the projection of the output gap on present, future, and past observations on output growth:

$$x_t = \sum_{j=-\infty}^{\infty} h_j \Delta y_{t-j} + \varepsilon_t^y, \tag{48}$$

where h_j is a scalar for each j and ε_t^y is uncorrelated with Δy_{t-s} for all s.[22] The projection that also involves unemployment can be expressed as follows:

$$x_t = \sum_{j=-\infty}^{\infty} h_j \Delta y_{t-j} + \sum_{j=-\infty}^{\infty} h_j^u u_{t-j} + \varepsilon_t^{y,u}.$$

Here, h_j^u is a scalar for each j and $\varepsilon_t^{y,u}$ is uncorrelated with Δy_{t-s}, u_{t-s} for all s. We define the information content of unemployment for the output gap by the ratio,

[21] That is, we set $\gamma = 1$ and $\psi = 0$.

[22] In practice only a finite amount of data is available. As a result, the projection involves a finite number of lags where the number of lags varies with t. The Kalman smoother solves the projection problem in this case.

$$r^{\text{two-sided}} \equiv \frac{E(\varepsilon_t^{y,u})^2}{E(\varepsilon_t^{y})^2}.$$

The lower the ratio, the greater the information in unemployment for the gap. We also compute the analogous variance ratio, $r^{\text{one-sided}}$, corresponding to the one-sided projection involving only current and past observations on the explanatory variables.[23] The one-sided projection is the one that is relevant to assess the information content of unemployment for policymakers working in real time. Our measure of information does not incorporate sampling uncertainty in parameters. The variances used to construct $r^{\text{two-sided}}$ and $r^{\text{one-sided}}$ assume the parameters are known with certainty and that the only uncertainty stems from the fact that the gap cannot be constructed using the data available to the econometrician.

3.3.2 The CTW model of unemployment

We convert the usual three equation log-linear representation of the New Keynesian model into a model of unemployment by adding one equation. This reduced-form log-linear system is derived from explicit microeconomic foundations in CTW. That paper also shows how our model of unemployment can be integrated into a medium-sized DSGE model such as the one in Section 4.

In the CTW model, finding a job requires exerting effort. Because effort only increases the probability of finding a job, not everyone who looks for a job actually finds one. The unemployed are people who look for a job without success. The unemployment rate is the number unemployed, expressed as fraction of the labor force. As in the official definition, the labor force is the number of people employed plus the number unemployed.

Since effort is unobserved and privately costly, perfect insurance against idiosyncratic labor market outcomes is not possible. As a result, the unemployed are worse off than the employed. In this way, the model captures a key reason that policymakers care about unemployment: a rise in unemployment imposes a direct welfare cost on the families involved. In this respect, our model differs from other work that integrates unemployment into monetary DSGE models.[24] In those models, individuals have perfect insurance against labor market outcomes.

We now describe the shocks and the linearized equilibrium conditions of the model. In previous sections of this chapter, the efficient level of hours worked is

[23] In the analysis below, we compute the projections in two ways. When we apply the filter to the data to extract a time series of x_t, we use the Kalman smoother. To compute the weights in the infinite projection problem, we use standard spectral methods described in, for example, Sargent (1979, Chapter 11). The spectral weights can also be computed by numerical differentiation of the output of the Kalman smoother with respect to the input data. We verified that the two methods produce the same results as long as the number of observations is large and t lies in the middle of the data set.

[24] For a long list of references, see CTW.

constant, and so the output gap can be expressed simply as the deviation of the number of people working from that constant (see Eq. 33). In this section, the efficient number of people working is stochastic. We denote the deviation of this number from steady state by h_t^*. We continue to assume that the steady state of our economy is efficient, so that $\hat{H}_t$ and h_t^* represent percent deviations from the same steady state values. The output gap is now:

$$x_t = \hat{H}_t - h_t^*.$$

The object, h_t^*, is driven by disturbances to the disutility of work, as well as by disturbances to the technology that converts household effort into a probability of finding a job. These various disturbances to the efficient level of employment cannot be disentangled using the data we assume are available to the econometrician. We refer to h_t^* as a labor supply shock. We hope that this label does not generate confusion. In our context this shock summarizes a broader set of disturbances than simply the one that shifts the disutility of labor. We adopt the following time series representation for the labor supply shock:

$$h_t^* = \lambda h_{t-1}^* + \varepsilon_t^{h^*}, \tag{49}$$

where $\varepsilon_t^{h^*}$ is a zero mean, iid process uncorrelated with h_{t-s}^*, $s > 0$ and $E\left(\varepsilon_t^{h^*}\right)^2 = \sigma_{h^*}^2$. In the version of the CTW model studied here, h_t^* is orthogonal to all the other shocks.

We assume the technology shock is a logarithmic random walk:

$$\Delta \log z_t = \varepsilon_t^z, \tag{50}$$

where Δ denotes the first difference operator. The object, ε_t^z, is a mean-zero, iid disturbance that is not correlated with $\log z_{t-s}$, $s > 0$. We denote its variance by $E\left(\varepsilon_t^z\right)^2 = \sigma_z^2$. The empirical rationale for the random walk assumption is discussed in Section 4.1.[25]

According to CTW, the interest rate in the first-best equilibrium is given by:

$$\hat{R}_t^* = E_t\left[\Delta \log z_{t+1} + h_{t+1}^* - h_t^*\right]. \tag{51}$$

Log consumption in the first best equilibrium is (apart from a constant term) the sum of $\log z_t$ and h_t^*. So, according to Eq. (51), $\hat{R}_t^*$ corresponds to the anticipated growth rate of (log) consumption. This reflects the CTW assumption that utility is additively

[25] Another way to assess the empirical basis for the random walk assumption exploits the simple model's implication that the technology shock can be measured using labor productivity. One measure of labor productivity is given by the ratio of real US GDP to a measure of total hours. The first-order autocorrelation of the quarterly logarithmic growth rate of this variable for the period, 1951Q1 to 2008Q4 is −0.02. The same first-order autocorrelation is 0.02 when calculated using output per hour for the nonfarm business sector. These results are consistent with our random walk assumption.

separable and logarithmic in consumption. We also suppose there is a disturbance, μ_t, that enters the Phillips curve as follows:

$$\hat{\pi}_t = \kappa^p x_t + \beta E_t \hat{\pi}_{t+1} + \mu_t, \tag{52}$$

where $\kappa^p > 0$. Here, κ^p denotes the slope of the Phillips curve in terms of the output gap. This is not to be confused with κ_p in Eqs. (41) and (35), which is the slope of the Phillips curve in terms of marginal cost. Our representation of the Phillips curve shock is given by

$$\mu_t = \chi \mu_{t-1} + \varepsilon_t^{\mu}, \tag{53}$$

where $E(\varepsilon_t^{\mu})^2 = \sigma_{\mu}^2$. The intertemporal equation, Eq. (42), is unchanged from before. Finally, we suppose that there is an iid disturbance, M_t, that enters the monetary policy rule in the following way:

$$\hat{R}_t = \rho_R \hat{R}_{t-1} + (1 - \rho_R)[r_\pi E_t \hat{\pi}_{t+1} + r_x x_t] + M_t, \tag{54}$$

where $E\left(\varepsilon_t^M\right)^2 = \sigma_M^2$. The four exogenous shocks in the model are orthogonal to each other at all leads and lags.

Let the unemployment gap, u_t^g, denote the deviation between actual unemployment and efficient unemployment, when both are expressed in percent deviation from their (common) nonstochastic steady state. The CTW model implies:

$$u_t^g = -\kappa^g x_t, \kappa^g > 0, \tag{55}$$

where κ^g is a function of underlying structural parameters. The previous expression resembles "Okun's law." If actual unemployment is one percentage point higher than its efficient level, then output is $1/\kappa^g$ percent below its efficient level. Discussions of Okun's law often suppose that $1/\kappa^g$ lies in a range of 2 to 3 (see, e.g., Abel & Bernanke, 2005). The unemployment rate in the efficient equilibrium, u_t^*, has the following representation:

$$u_t^* = -\omega h_t^*, \omega > 0.$$

In the CTW model, the factors that increase labor supply also increase the intensity of job finding effort, and this is the reason unemployment in the efficient equilibrium falls. The harder people look for a job, the sooner they find what they are looking for. According to the previous two equations, the actual unemployment rate, u_t, satisfies the following equation:

$$u_t = -\left[\kappa^g x_t + \omega h_t^*\right]. \tag{56}$$

Absent the presence of the labor supply shock, the efficient level of unemployment would be constant and the actual unemployment rate would represent a direct observation on the output gap.

In sum, the log-linearized equations of the CTW model consist of the usual three equations of the standard New Keynesian model, Eqs. (42), (52), and (54), plus an equation that characterizes unemployment, Eq. (56). In addition, there are the equations that characterize the laws of motion of the exogenous shocks and of the efficient rate of interest.

3.3.3 Limited information Bayesian inference

To investigate the quantitative implications of the model, we must assign values to its parameters. We set values of the economic parameters of the model,

$$\kappa^p, \omega, \kappa^g, r_\pi, r_x, \rho_R, \beta,$$

as indicated in Table 1a. Let θ denote the 6×1 column vector consisting of the parameters governing the stochastic processes:

$$\theta = (\lambda, \chi, \sigma_z, \sigma_{h^*}, \sigma_M, \sigma_\mu)'. \tag{57}$$

We use data on output growth and unemployment to select values for the elements of θ.[26] We do this using a version of the limited information Bayesian procedure described in Kim (2002) and in Section 5.2.[27] Let γ denote the 11×1 column vector composed of the j^{th} order autocovariance matrix of output growth and unemployment, for $j = 0, 1, 2$. Let $\hat{\gamma}$ denote the corresponding sample estimate based on $T = 232$ quarterly observations, 1951Q1–2008Q4. Hansen's (1982) generalized method of moments analysis (GMM) establishes that for sufficiently large T, $\hat{\gamma}$ is a realization from a Normal distribution with mean equal to the true value of the second moments, γ^0, and variance, V/T. The results also hold when V is replaced by a consistent sample estimate, $\hat{V}$.[28] Our model provides a mapping from θ to γ, which we denote by $\gamma(\theta)$. Hansen's

[26] The seasonally adjusted unemployment rate for people 16 years and older was obtained from the Bureau of Labor Statistics and has mnemonic LNS14000000. We use standard real per capita GDP data, as described in Section A of the technical appendix found at http://www.faculty.econ.northwestern.edu/faculty/christiano/research/Handbook/technical_appendix.pdf.

[27] The procedure is the Bayesian analog of the moment matching estimation procedure described in Christiano and Eichenbaum (1992a).

[28] We compute $\hat{V}$ as follows. Let θ^0 denote the true, but unknown, values of the model parameters. Let $h(\gamma, w_t)$ denote the 11×1 GMM error vector having the property, $Eh(\gamma^0, w_t) = 0$, where $w_t = (\Delta y_t, u_t)'$. Let $g_T(\gamma) \equiv (1/T)\sum_t h(\gamma, w_t)$ and define $\hat{\gamma}$ by $g_T(\hat{\gamma}) = 0$. Then, $\sqrt{T}(\hat{\gamma} - \gamma^0)$ converges in distribution as $T \to \infty$ to $N(0, V)$. Here, $V = (D')^{-1} SD^{-1}$, where S denotes the spectral density at frequency zero of $h(\gamma^0, w_t)$, and $D' = \lim_{T\to\infty} \partial g_T(\gamma)/\partial \gamma'$, where the derivative is evaluated at $\gamma = \hat{\gamma}$ (for a discussion of these convergence results of GMM see, e.g., Hamilton 1994). Our estimator of V is given by $\hat{V} = (\hat{D}')^{-1} \hat{S}\hat{D}^{-1}$. We estimate S by $\hat{\Gamma}_0 + (1 - 1/3)\left(\hat{\Gamma}_1 + \hat{\Gamma}'_1\right) + (1 - 2/3)\left(\hat{\Gamma}_2 + \hat{\Gamma}'_2\right)$, where $\hat{\Gamma}_j = [\sum_t h(\hat{\gamma}, w_t)\, h(\hat{\gamma}, w_{t-j})']/(T - j)$, $j = 0, 1, 2$. Also, $\hat{D}$ is D with unknown true parameters replaced by consistent estimates. An alternative version of our limited information Bayesian strategy, which we did not explore, works with $V(\theta)$, which is the V matrix constructed with the D and S matrices implied by the model when its parameter values are set to θ.

Table 1a Non-estimated Parameters in Simple Model

Parameter	Value	Description
β	0.99	Discount factor
r_π	1.5	Taylor rule: inflation coefficient
r_x	0.2	Taylor rule: output gap coefficient
ρ_R	0.8	Taylor rule: interest rate smoothing coefficient
κ^p	0.11	Slope of Phillips curve
κ^g	0.4	Okun's law coefficient
ω	1.0	Elasticity of efficient unemployment, u^*, w.r.t. efficient hours, h^*

result suggests that, for sufficiently large T, the likelihood of $\hat{\gamma}$ conditional on θ and $\hat{V}$ is given by the following multivariate Normal distribution:[29]

$$p\left(\hat{\gamma}|\theta;\frac{\hat{V}}{T}\right)=\frac{1}{(2\pi)^{6/2}}\left|\frac{\hat{V}}{T}\right|^{-\frac{1}{2}}\exp\left\{-\frac{T}{2}(\hat{\gamma}-\gamma(\theta))'\hat{V}^{-1}(\hat{\gamma}-\gamma(\theta))\right\}. \tag{58}$$

Given a set of priors for θ, $p(\theta)$, the posterior distribution of θ conditional on $\hat{\gamma}$ and $\hat{V}$ is, for sufficiently large T,

$$p\left(\theta|\hat{\gamma};\frac{\hat{V}}{T}\right)=\frac{p\left(\hat{\gamma}|\theta;\frac{\hat{V}}{T}\right)p(\theta)}{p\left(\hat{\gamma};\frac{\hat{V}}{T}\right)}.$$

[29] We performed a small Monte Carlo experiment to investigate whether Hansen's asymptotic results are likely to be a good approximation with a sample size, $T=232$. The results of the experiment make us cautiously optimistic. Our Monte Carlo study used the classic New Keynesian model without unemployment (i.e., Eqs. (42), (50), and (51) with $h_t^*\equiv 0$, and Eqs. (52)–(54). With one exception, we set the relevant economic parameters as in Table 1a. The exception is ρ_R, which we set to zero. In addition, the parameters in θ were set as in the posterior mode for the partial information procedure in Table 1b. With this parameterization, the model implies (after rounding) $\sigma_y=0.021$, $\rho_1=\rho_2=-0.039$. Here, $\sigma_y\equiv E(\Delta y_t)^2]^{1/2}$, $\rho_i\equiv E(\Delta y_t\Delta y_{t-i})/\sigma_y^2$, $i=1,2$. We then simulated 10,000 data sets, each with $T=232$ artificial observations on output growth, Δy_t. The mean, across simulated samples, of estimates of σ_y, ρ_1, ρ_2, is, respectively, 0.021, −0.039, and −0.033. Thus, the results are consistent with the notion that our second moment estimator is essentially unbiased. To investigate the accuracy of Hansen's Normality result, we examined the coverage of 80% confidence intervals computed in the usual way (i.e., the point estimate ±1.28 times the corresponding sample standard deviation computed in exactly the way specified in the previous footnote). In the case of σ_y, ρ_1, ρ_2 the 80% confidence interval excluded the true values of the parameters 22.35, 21.87, and 21.39% of the time, respectively. We found these to be reasonably close to the 20% numbers suggested by the asymptotic theory. Related to this, we found little bias in our estimator of the sample standard deviation estimator. In particular, the actual standard deviation of the estimator of σ_y, ρ_1, and ρ_2 across the 10,000 samples is 0.00098, 0.064, and 0.065. The mean of the corresponding sample estimates is 0.00095, 0.062, and 0.064, respectively. Evidently, the estimator of the sampling standard deviation is roughly unbiased.

The marginal density, $p(\hat{\gamma};\ \hat{V}/T)$, as well as the marginal posterior distribution of individual elements of θ can be computed using a standard random walk metropolis algorithm or by using the Laplace approximation.[30] In the present application, we use the Laplace approximation. Our moment-matching Bayesian approach has several attractive features. First, it has the advantage of transparency because it focuses on a small number of features of the data. Second, it does not require the assumption that the underlying data are realizations from a Normal distribution, as is the case in conventional Bayesian analyses.[31] The Normality in Eq. (58) depends on the validity of the central limit theorem, not on Normality of the underlying data. Third, the method has the advantage of computational speed. The matrix inversion and log determinant in Eq. (58) needs to be computed only once. In addition, evaluating a quadratic form like the one in Eq. (58) is computationally very efficient. These computational advantages are likely to be important when searching for the mode of the posterior distribution. Moreover, the advantages may be overwhelmingly important when computing the whole posterior distribution using a standard random walk Metropolis algorithm. In this case, Eq. (58) must be evaluated on the order of hundreds of thousands of times.

Because our econometric method may be of independent interest, we compare the results obtained using it with results based on a conventional full information Bayesian approach. In particular, let Y denote the data on unemployment and output growth used to compute $\hat{\gamma}$ for our limited information Bayesian procedure. In this case, the posterior distribution of θ given Y is

$$p(\theta|Y) = \frac{p(Y|\theta)p(\theta)}{p(Y)},$$

where $p(Y|\theta)$ is the Normal likelihood function and $p(Y)$ is the marginal density of the data. The priors, $p(\theta)$, used in the two econometric procedures are the same and they are listed in Table 1b.

Table 1b reports posterior modes and posterior standard deviations for the parameters, θ. Note how similar the results are between the full and limited information methods. The one difference has to do with λ, the autoregressive parameter for the labor supply shock. The posterior mode for this parameter is somewhat sensitive to which econometric method is used. The standard deviation of the posterior mode of λ is more sensitive to the method used. In all but one case, there appears to be substantial information in the data about the parameters, as measured by the reduction in standard deviation from prior to posterior. The exception is λ. Under

[30] For additional discussion of the Laplace approximation, see Section 5.4.
[31] Failure of Normality in aggregate macroeconomic data is discussed in Christiano (2007).

Table 1b Priors and Posteriors for Parameters of Simple Model

		Prior		Posterior mode [Std. dev.[a]]	
Parameter		**Distribution [bounds]**	**Mean, std. dev. [5% and 95%]**	**Limited info[b]**	**Full info[c]**
Exogenous processes parameters					
Autocorrelation, labor supply shock	λ	Beta	0.75, 0.15	0.71	0.83
		[0, 1]	[0.47, 0.95]	[0.16]	[0.08]
Autocorrelation, Phillips curve shock	χ	Beta	0.75, 0.15	0.92	0.93
		[0, 1]	[0.47, 0.95]	[0.01]	[0.02]
Std. dev., technology shock (%)	σ_z	Inv. gamma	0.50, 0.40	0.62	0.63
		[0, ∞]	[0.18, 1.04]	[0.04]	[0.04]
Std. dev., labor supply shock (%)	σ_{h^*}	Inv. gamma	0.50, 0.40	0.24	0.19
		[0, ∞]	[0.18, 1.04]	[0.06]	[0.03]
Std. dev., monetary policy shock (%)	σ_M	Inv. gamma	0.50, 0.40	0.13	0.11
		[0, ∞]	[0.18, 1.04]	[0.01]	[0.01]
Std. dev., Phillips curve shock (%)	σ_μ	Inv. gamma	0.50, 0.40	0.24	0.25
		[0, ∞]	[0.18, 1.04]	[0.03]	[0.03]

[a]Based on Laplace approximation.
[b]Limited info refers to our Bayesian moment–matching procedure.
[c]Full info refers to standard full information Bayesian inference based on the full likelihood of the data.

the limited information procedure, there is little information in the data about this parameter.

We analyze the properties of the model at the mode of the posteriors of θ. Because the Table 1b results are so similar between limited and full information methods, the corresponding model properties are also essentially the same. As a result, we only report properties based on the posterior mode implied by the limited information procedure.

Table 1c reports $\hat{\gamma}$, the empirical second moments underlying the limited information estimator, as well as the corresponding second moments implied by the model. The empirical and model moments are reasonably close. The variance decomposition implied by the model is reported in Table 1d. Most of the variance in output is due to technology shocks and to the disturbance in the Phillips curve. Note that technology shocks have no impact on any of the other variables. This reflects that with our policy rule, the economy's response to a random walk technology shock is efficient and involves no response in the interest rate, inflation, or any labor market variable. The economics of this result is discussed in Section 3.4. In the case of unemployment, the disturbance to the Phillips curve is the principle source of fluctuations. Labor supply shocks turn out to be relatively unimportant as a source of fluctuations. The implications of the latter finding for our results are discussed in the next section.

3.3.4 Estimating the output gap using the CTW model

The implications of our model for the information in the unemployment rate for the output gap is displayed in Table 1e. The row called "posterior mode" reports

$$r^{\text{two-sided}} = 0.11 \text{ and } r^{\text{one-sided}} = 0.09.$$

Table 1c Properties of Simple Model (at Limited Information Posterior Mode) and Data[a]

Covariances (× 100)	Model	Data	Covariances (× 100)	Model	Data
Cov. $(\Delta y_t, \Delta y_t)$	0.0099	0.0090	Cov. $(\Delta y_t, \Delta y_{t-2})$	0.0010	0.0017
Cov. (u_t, u_t)	0.0190	0.0220	Cov. $(\Delta y_t, u_{t-2})$	0.0021	0.0033
Cov. $(\Delta y_t, u_t)$	−0.0013	−0.0002	Cov. $(u_t, \Delta y_{t-2})$	−0.0025	−0.0038
			Cov. (u_t, u_{t-2})	0.0174	0.0201
Cov. $(\Delta y_t, \Delta y_{t-1})$	0.0021	0.0030			
Cov. $(\Delta y_t, u_{t-1})$	0.0012	0.0022			
Cov. $(u_t, \Delta y_{t-1})$	−0.0021	−0.0023			
Cov. (u_t, u_{t-1})	0.0184	0.0215			

[a]Sample: 1951Q1 to 2008Q4. Data series: Δy — real per capita GDP growth, u — unemployment rate.

Table 1d Variance Decomposition of Simple Model (at Limited Information Posterior Mode, in %)

Output growth	Unemployment rate	Nom. interest rate	Inflation rate	Output gap
		Technology Shocks		
38.7	0	0	0	0
		Monetary Policy Shocks		
17.7	1.8	0.7	0.5	1.9
		Labor Supply Shocks		
0.7	3.9	0.1	0	0.3
		Phillips Curve Shocks		
42.9	94.3	99.2	99.5	97.8

Table 1e Information About Output Gap in Unemployment Rate, *u*, Simple Model

	Two-sided projection			**One-sided projection**		
	Projection error (%)			**Projection error (%)**		
	standard deviation×100			**standard deviation×100**		
Parameter	***u*** **Observed**	***u*** **Unobserved**	$r^{\text{two-sided}}$	***u*** **Observed**	***u*** **Unobserved**	$r^{\text{one-sided}}$
Posterior mode	0.74	2.26	*0.11*	0.79	2.66	*0.09*
Alternative parameter values						
$\lambda = 0.99999$, $100\sigma_{h^*} = 0.0015$	0.68	2.24	*0.09*	0.68	2.64	*0.07*
$\omega = 0.001$	0.00081	2.26	*0.00*	0.00084	2.65	*0.00*
$100\,\sigma_{h^*} = 0.001$	0.0036	2.24	*0.00*	0.0036	2.64	*0.00*
$100\,\sigma_{h^*} = 1$	1.80	2.53	*0.51*	2.12	2.84	*0.56*

Note: (i) $r^{\text{two-sided}}$ is the ratio of the two-sided projection error variance when *u* is observed to what it is when it is not observed. $r^{\text{one-sided}}$ is the analogous object for the case of one-sided projections. For details, see the text. (ii) The posterior mode of the parameters are based on our limited information Bayesian procedure.

Thus, in the case of the two-sided projection, the variance of the projection error in the output gap is reduced by 89% when the unemployment rate is included in the data used to estimate the output gap. The 95% confidence interval for the percent output gap is the point estimate plus and minus 4.4% when the estimate is based only on

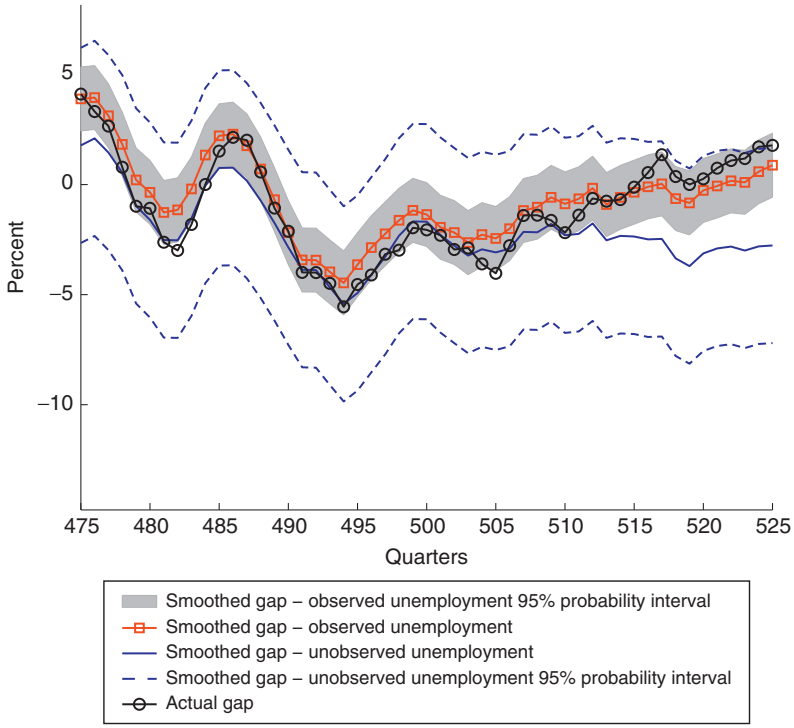

Figure 2 Actual versus smoothed output gap, artificial data.

the output growth data. That interval shrinks by over 60%, to ±1.5% with the introduction of unemployment.[32] Figure 2 displays observations 475 to 525 in a simulation of 1000 observations from our model. The figure shows the actual gap as well as estimates-based information sets that include only output growth and output growth plus unemployment. In addition, we display 95% confidence tunnels corresponding to the two information sets.[33] Note how much wider the tunnel is for estimates based on output growth alone.

Our optimal linear estimator of the output gap based on output growth alone (see Eq. 48) is directly comparable to the HP filter as an estimator of the gap.[34] The latter is also based on output data alone. The information in Figure 3 allows us to compare these two filters. Panel a shows the filter weights as they apply to the level of output,

[32] These observations are based on the following calculations: $1.5 = 0.0074 \times 1.96 \times 100$ and $4.4 = 0.0226 \times 1.96 \times 100$ using the information in Table 1e. Here, 1.96 is the 2.5% critical value for the standard Normal distribution.

[33] The confidence tunnels are constructed by adding and subtracting 1.96 times the standard deviation of the projection error standard deviation implied by the Kalman smoother to the smoothed estimates of the gap. The assumption of Normality implicit in multiplying by 1.96 is justified here because the disturbances in the underlying simulation are drawn from a Normal distribution.

[34] We set the smoothing parameter in the HP filter to 1600.

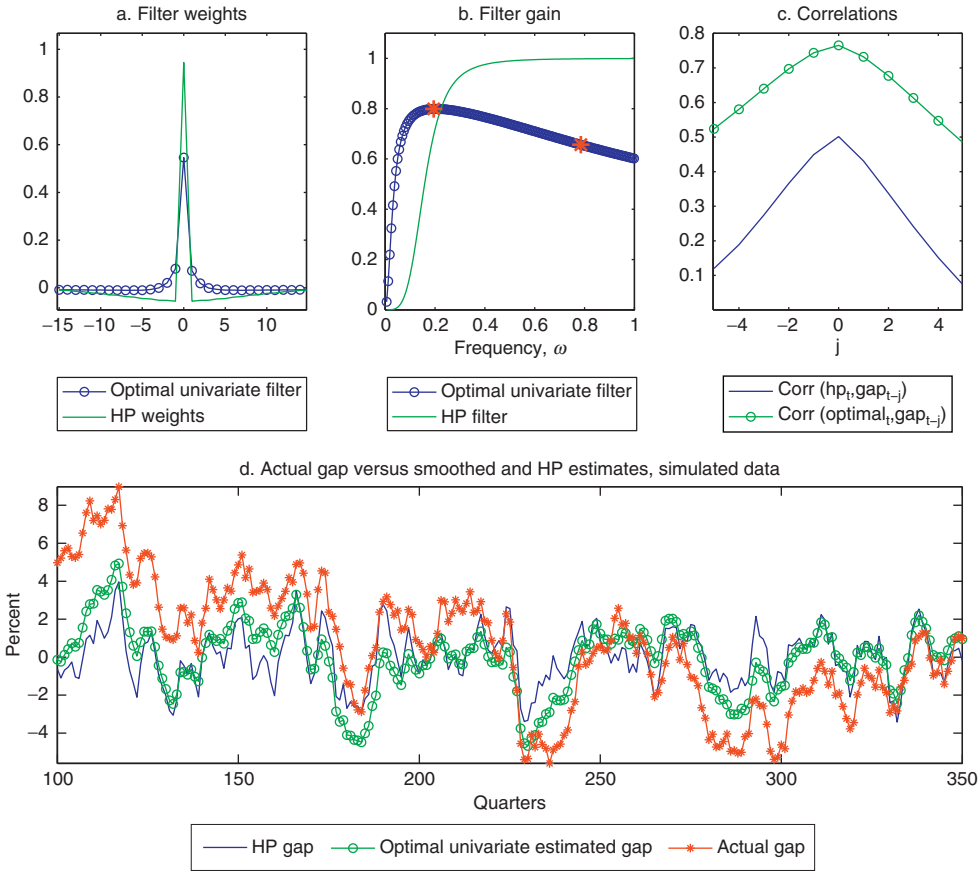

Figure 3 HP Filter and optimal univariate filter for estimating output gap. Note: Stars in panel b indicate business cycle frequencies corresponding to 2 and 8 years.

γ_t.[35] Note how similar the pattern of weights is, although they are not identical. The filter weights for the HP filter are known to be exactly symmetric. This is not a property of the optimal weights. However, panel a in Figure 3 shows that the optimal filter weights are very nearly symmetric. So, while the phase angle of the HP filter is exactly zero, the phase angle of the optimal filter implied by our model is nearly zero. Panel b in Figure 3 compares the gain of the two filters over a subset of frequencies that includes the business cycle frequencies, whose boundaries are indicated in the figure by stars. Evidently, both are approximately high pass filters. However, the optimal filter

[35] We computed the filter weights for the HP filter as well as for Eq. (48) by expressing the filters in the frequency domain and applying the inverse Fourier transform. In the case of Eq. (48), we compute the $\tilde{h}_j$'s in $x_t = \sum_{j=-\infty}^{\infty} h_j \Delta y_{t-j} + \varepsilon_t^y = \sum_{j=-\infty}^{\infty} \tilde{h}_j y_{t-j} + \varepsilon_t^y$.We use the result in King and Rebelo (1993) to express the HP filter in the frequency domain.

lets through lower frequency components of the output data and also slightly attenuates the higher frequencies. Panel c displays the cross correlations of the actual output gap with the HP and optimal filters, respectively. This was done in a sample of 1000 artificial observations on output simulated from our model (the optimal filter in a finite sample of data is obtained using the Kalman smoother). Note that both estimates are positively correlated with the actual gap. Of course, the gap is more highly correlated with the optimal estimate of that gap than with the HP filter estimate. Panel d in Figure 3 displays a subsample of our artificial data. We can see directly how similar the two filters are. However, note that there is a substantial low frequency component in the actual gap and this low frequency component is better tracked by the optimal filter. This is consistent with the result in panel b which indicates that the optimal filter allows lower frequency components of output to pass through.

Next, we applied the same statistical procedure to the U.S. data that we used to estimate the output gap in the artificial data. The results are displayed in Figure 4, which displays HP-filtered, log, real, per capita Gross Domestic Product (GDP), as well as the two-sided estimate of the output gap when unemployment is and is not included in the data set used in the projections.[36] We have not included confidence tunnels, to

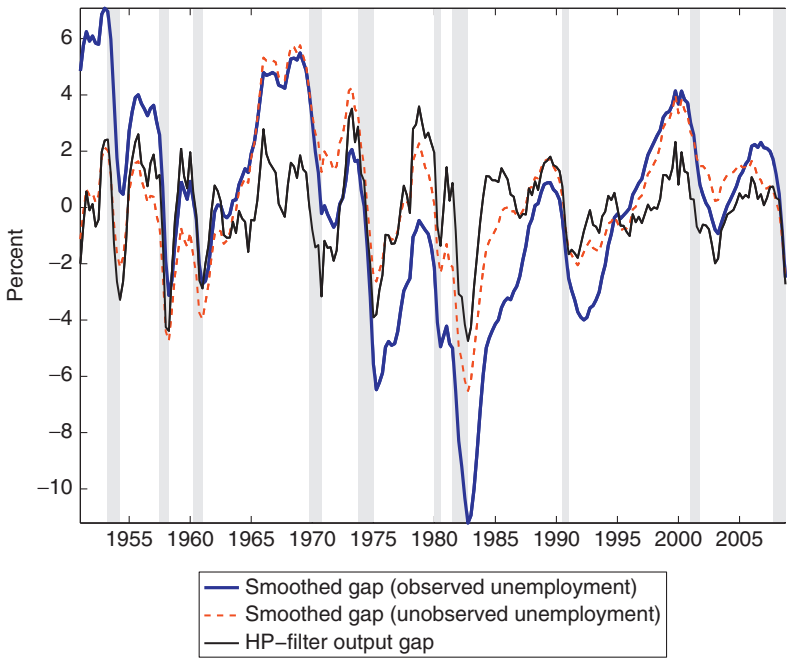

Figure 4 Output gap in U.S. data.

[36] Our calculations for Figure 4 are based on the Kalman smoother.

avoid further cluttering the diagram. In addition, the gray areas in the bracket denote the start and end date of recessions, according to the National Bureau of Economic Research (NBER). Several observations are worth making about the results in Figure 4. First, the estimated output gap is always relatively low in a neighborhood of NBER recessions. Second, the gap shows a tendency to begin falling before the onset of the NBER recession. This is to be expected. The NBER typically dates the start of a recession by the first quarter in which the economy undergoes two quarters of negative growth. Given that growth in the U.S. economy is positive on average, the start date of the NBER recession occurs after economic activity has already been winding down for at least a few quarters. This also explains why the HP filter estimate of the gap also typically starts to fall one or two quarters before an NBER recession. Third, consistent with the results in the previous paragraph, the gap estimates based on the HP filter and our estimate based on output data alone produce very similar results. Fourth, the inclusion of unemployment in the data used to estimate the output gap has a quantitatively large impact on the results. The estimated gap is substantially more volatile when unemployment is used and it is also more volatile than the HP filter gap. That the incorporation of unemployment has a big impact is perhaps not surprising, given the posterior mode of our parameters, which implies that labor supply shocks are relatively unimportant. As a result, the efficient unemployment rate, u_t^*, is not very volatile and the actual unemployment rate is a good indicator of the output gap (see Eq. 55).

We gain additional insight into our measures of the gap by examining the implied estimates of potential output. These are presented in Figure 5, which displays actual output as well as our measures of potential output based on using just output and using output and unemployment. Not surprisingly, in view of the results in Figure 4, the estimate of potential that uses unemployment is the smoother one of the two. Our results are similar to the results presented by Justiniano and Primiceri (2008), who also conclude that potential output is smooth.[37]

Our model is well suited to shed light on the question: Under what circumstances can we expect unemployment to contain useful information about the output gap? The general answer is that if the efficient level of unemployment is constant, then the actual unemployment rate is highly informative, because in this case it represents a direct observation on the output gap. This is documented in three ways in Table 1e. First, we consider the case where the total variance in the labor supply shock, h_t^*, is kept constant, but is reallocated into the very low frequencies. A motivation for this is the

[37] Although Sala, Söderström and Trigari (2008) do not specifically display their model's implications for potential output, one can infer from their estimate of the output gap that the measure of potential output implicit in their calculations is also smooth, like the one presented by Justiniano and Primiceri (2008). Except for these two papers, estimates of potential GDP reported in the literature are often more volatile than what we find. See, for example, Walsh's (2005) discussion of Levin, Onatski, Williams, and Williams (2005). See also Kiley (2010) and the sources he cites.

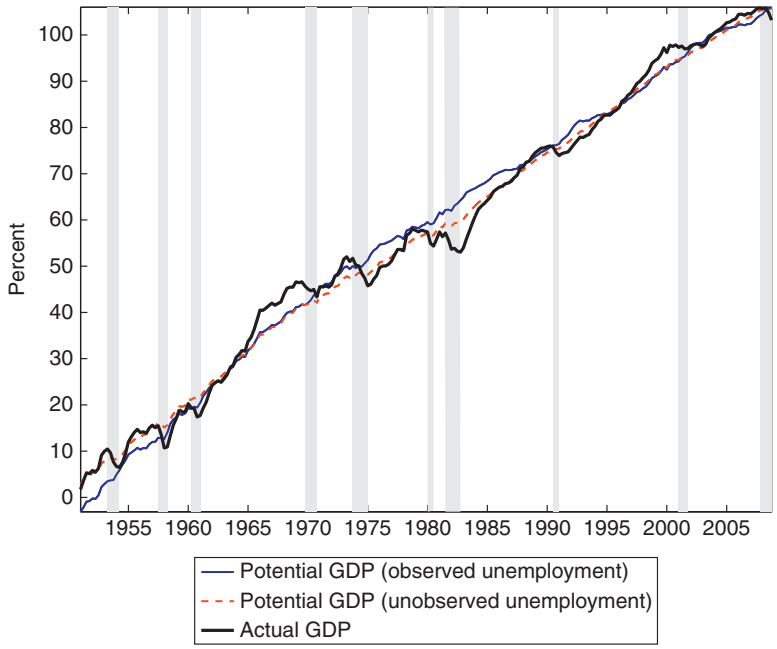

Figure 5 Actual output and two measures of potential output, U.S. data.

finding in Christiano (1988, pp. 266–268) that a low frequency labor supply shock is required to accommodate the behavior of aggregate hours worked. We set $\lambda = 0.99999$ and adjust $\sigma^2_{h^*}$ so that the variance of h^*_t is equal to what is implied by the model at the posterior mode. In this case, the efficient level of employment is a variable that evolves slowly over time.[38] As a result, the efficient rate of unemployment itself is slow-moving, so that most of the short-term fluctuations in the actual unemployment rate correspond to movements in the unemployment gap, u^g_t, and, hence in the output gap (recall Eq. 55). Consistent with this intuition, Table 1e indicates that the increase in λ causes $r^{\text{two-sided}}$ and $r^{\text{one-sided}}$ to fall to 0.09 and 0.07, respectively. Similarly, Table 1e also shows that if we reduce the magnitude of ω or of the variance of the labor supply shock itself, then the use of unemployment data essentially removes all uncertainty about the output gap. Finally, the table also shows what happens when we increase the importance of the labor supply shock. In particular, we increased the innovation variance in h^*_t by a factor of 4, from 0.24% to 1.0%. The result of this change on the model is that labor supply shocks now account for 10% of the variance of output growth and 41% of the variance of unemployment. With the efficient level of unemployment more volatile, we can expect that the value of the unemployment

[38] This captures the view that the evolution of h^*_t represents demographic and other slowly-moving factors.

rate for estimating the output gap is reduced. Interestingly, according to Table 1e, unemployment is still very informative for the output gap. Despite the relatively high volatility in the labor supply shock, the unemployment rate still reduces the variance of the prediction error for the output gap by about 44–49%.

In sum, the results reported here suggest the possibility that the unemployment rate might be very useful for estimating the output gap. We find that this is likely to be particularly true if the efficient level of unemployment evolves slowly over time. In addition, we found in our estimated model that the HP filter estimate of the gap closely resembles the estimate of the gap that is optimal conditional on our model. All of these observations ought to be viewed as suggestive at best. Because part of our objective here is pedagogic, the observations were made in a very simple setting. It would be interesting to investigate whether they are also true in more complicated environments with more shocks, in which more data are available to the econometrician. Section 3.4 shows that the optimal filter for extracting the output gap is very sensitive to the details of the underlying model. As a consequence, the similarity between the HP filter and the optimal filter found in this section ought to only be treated as suggestive. A final assessment of the relationship between the two filters requires additional experience with a variety of models.

3.4 Using HP-filtered output to estimate the output gap

Section 3.3.4 displayed a model environment with the property that the HP filter is nearly optimal as a device for estimating the output gap. This section shows that the accuracy of the HP filter for extracting the output gap is very sensitive to the details about the underlying model. We demonstrate this point in a simple version of the classic New Keynesian model (i.e., $\gamma = 1$, $\psi = 0$) in which there is only one shock, the technology shock. We show that the HP filter may be positively or negatively correlated with the true output gap, depending on the time series properties of the shock. When the shock triggers strong wealth effects, then output overreacts to the shock, relative to the efficient equilibrium. In this case, the HP-filtered estimate of the gap is positively correlated with the true output gap. If the shock triggers only a weak wealth effect, that correlation is negative.

Our analysis requires a careful review of the economics of the response of employment and output to a technology shock. This is a topic that is of independent interest because it has attracted widespread attention, primarily in response to the provocative paper by Gali (1999).

The linearized equilibrium conditions of the model are given by Eqs. (40)–(43), with $\psi = 0$, $\gamma = 1$. We consider the following two laws of motion for technology:

$$\Delta \log z_t = \rho_z \Delta \log z_{t-1} + \varepsilon_t^z, \text{“AR (1) in growth rate”}$$
$$\log z_t = \rho_z \log z_{t-1} + \varepsilon_t^z, \text{“AR (1) in levels”}.$$

These two laws of motion have the same implication for what happens to z_t in the period of a positive realization of ε_t^z. But, they differ sharply in their implications for the eventual impact of a shock on z_t. In the AR(1) in growth rate, a 0.01 shock in ε_t^z drives up z_t by 1%, but creates the expectation that z_t will eventually rise by $1/(1-\rho_z)$%. In the AR(1) in levels representation, a jump in z_t is associated with the expectation that z_t will be lower in later periods. We adopt the following parameterization:

$$\beta = 0.99, \rho_z = 0.5, \rho_R = 0, r_x = 0.2, r_\pi = 1.5, \phi = 0.2, \xi_p = 0.75.$$

In the case of the AR(1) in growth rate, a 1% shock up in technology is followed by additional increases, with technology eventually settling at a level that is permanently higher by 2% (see panel d in Figure 6). The response of the efficient level of consumption coincides with the response of the technology shock. Households in this economy experience a big rise in wealth in the moment of the shock. The motive to smooth consumption intertemporally makes them want to set their consumption to its permanently higher level right away. The rise in the rate of interest in the efficient

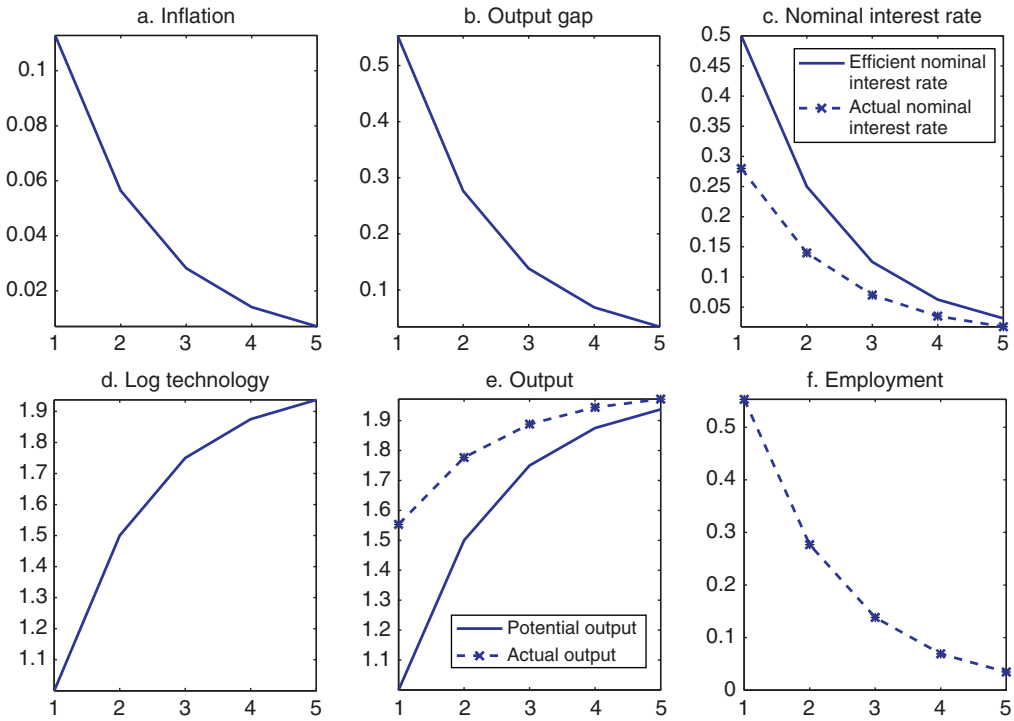

Figure 6 Dynamic response of simple model without capital to a one percent technology shock (AR(1) in growth rate specification).

equilibrium is designed to restrain this potential surge in consumption. This is why it is that in the efficient equilibrium, output (see panel e Figure 6) rises by the same amount as the technology shock, while employment remains unchanged. Now consider the actual equilibrium. According to panel c in the figure, the interest rate rule generates an inefficiently small rise in the rate of interest. As a result, monetary policy fails to fully reign in the surge in consumption demand triggered by the shock. Employment rises and so output rises by more than the technology shock. The increase in employment leads to an increase in costs and, therefore, inflation. The output gap responds positively to the shock and so potential output (i.e., the efficient level of output) is less volatile than the actual level. We can expect that the output gap estimated by the HP filter, which estimates potential output smoothing actual output, will at least be positively correlated with the true output gap.

We simulated a large number of artificial observations using the model and we then HP-filtered the output data.[39] Figure 7A displays actual, potential and HP smoothed

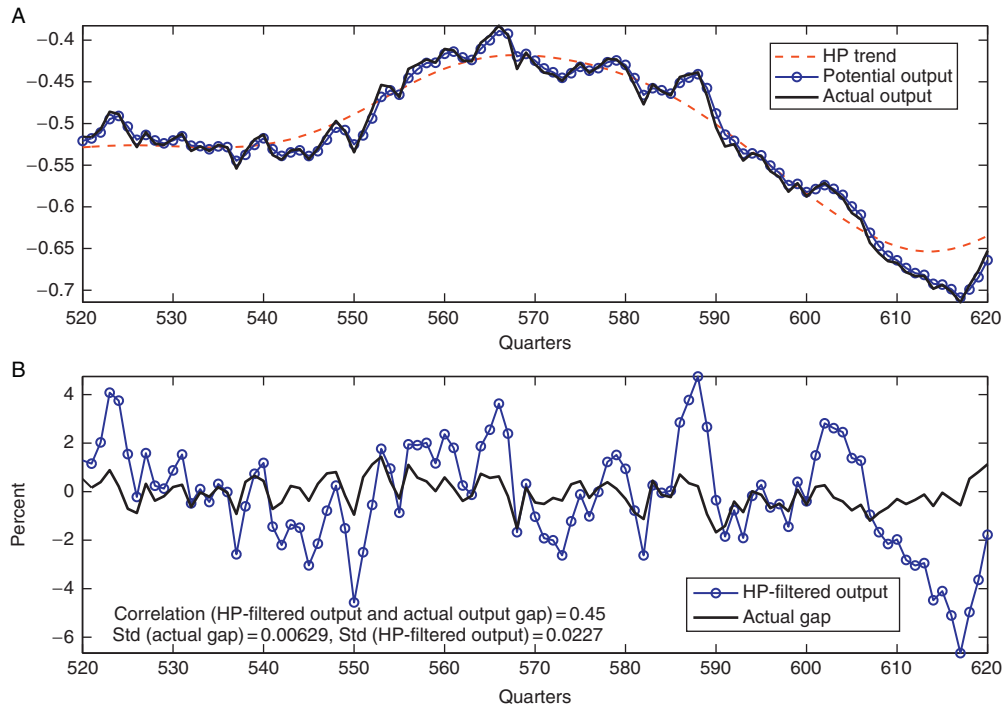

Figure 7 (A) Potential output, actual output and hp trend based on actual output (simulated data) (B) HP filter estimate of output gap versus actual gap (simulated data). AR(1) in growth rate specification.

[39] We used the usual smoothing parameter value for quarterly data, 1600.

output. We can see that the HP filter substantially oversmooths the data. However, consistent with the presumption implicit in the HP filter, the actual level of output is (somewhat) more volatile than the corresponding efficient level. Figure 7B displays the actual gap and the HP-estimated gap. Note that they are positively correlated, though the HP filtered gap is too volatile.

Now consider the AR(1) in levels specification of technology. The dynamic response of technology to a 1% disturbance in ε_t^z is displayed in panel d of Figure 8. The state of technology is high in the period of the shock, compared to its level anticipated for later periods. As before, the efficient level of consumption mirrors the time path of the technology shock. In the efficient equilibrium, agents expect lower future consumption and so intertemporal smoothing motivates them to cut current consumption relative to its efficient level. The drop in the interest rate in the efficient equilibrium is designed to resist this relative weakness in consumption (see panel c). Put differently, a sharp drop in the interest rate is needed to ensure that demand expands by enough to keep employment unchanged in the face of the technology improvement. In the actual equilibrium, the monetary policy rule cuts the interest rate less aggressively than in the

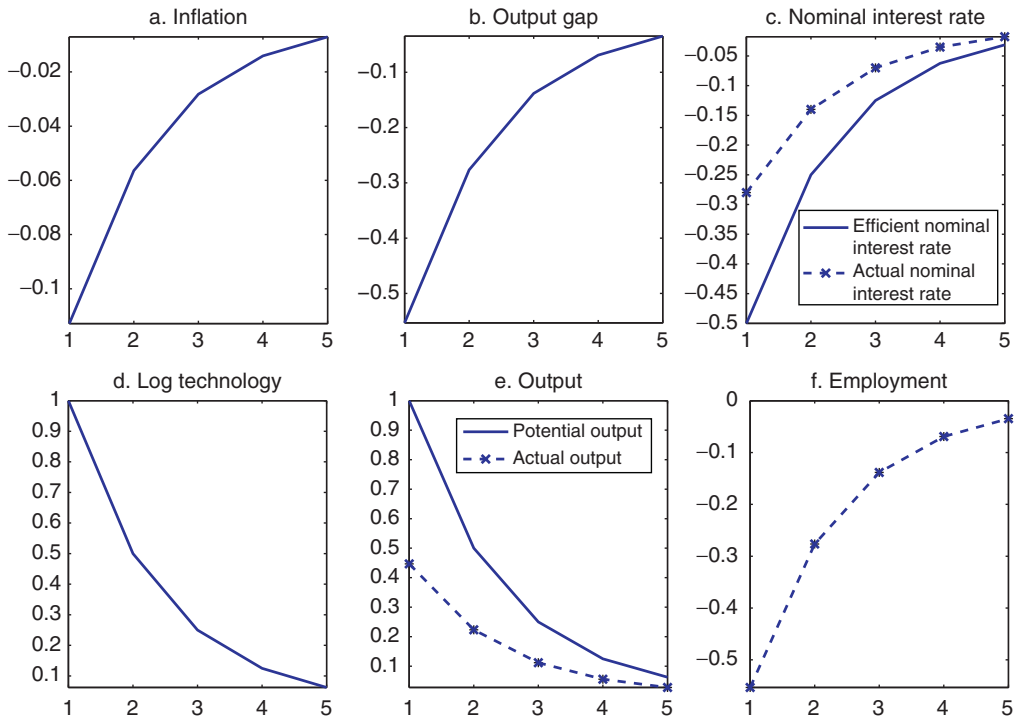

Figure 8 Dynamic response of simple model without capital to a one percent technology shock. AR(1) in levels specification.

efficient equilibrium. The relatively small drop in the interest rate fails to reverse the weakness in demand. As a result, the response of output is relatively weak and employment falls. The fall in employment is associated with a fall in marginal production costs and this explains why inflation falls in response to the technology shock. Figure 9A displays the implications of the AR(1) in levels specification of technology for the HP filter as a way to estimate the output gap. Note how potential output is substantially more volatile than actual output. As an estimator of potential output, the HP filter goes in precisely the wrong direction, by smoothing. Figure 9B compares the HP filter estimate of the output gap with the corresponding actual value. Note how the two are now negatively correlated.

A by-product of the previous discussion is an exploration of the economics of the response of hours worked to a technology shock in the classic New Keynesian model. In that model, hours worked rise in response to a technology shock that triggers a big wealth effect, and falls in response to a technology shock that implies a weak wealth effect. The principle that the hours worked response is greater when a technology shock triggers a large wealth effect survives in more complicated New Keynesian models such as the one discussed in the next section.

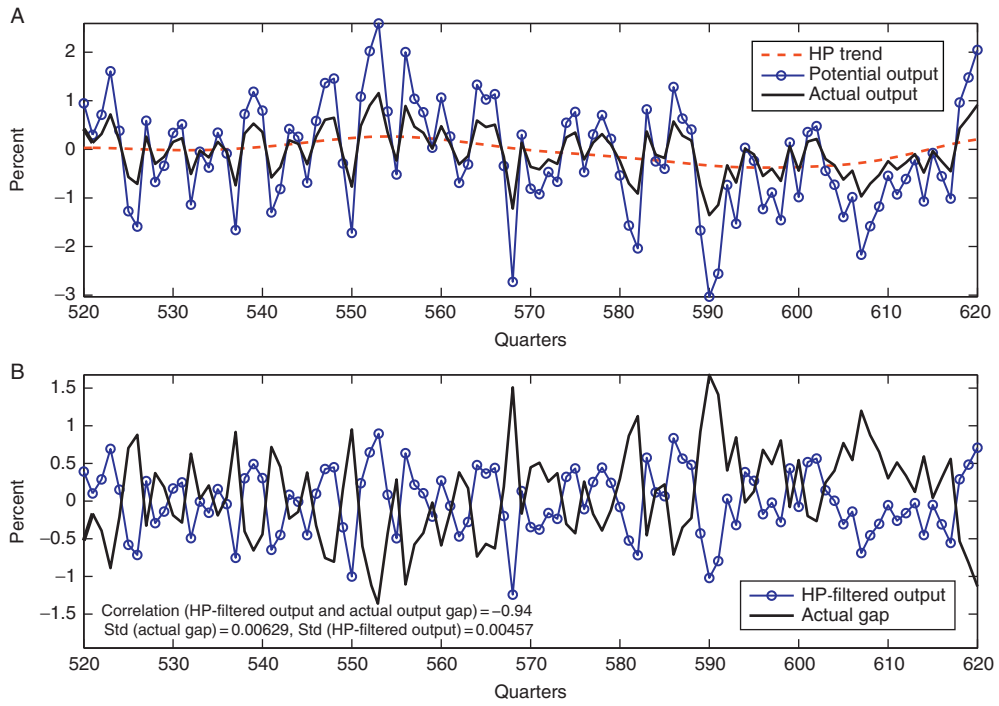

Figure 9 (A) Potential output, actual output and hp trend based on actual output (simulated data) (B) HP filter estimate of output gap versus actual gap (simulated data). AR(1) in levels specification.

4. MEDIUM-SIZED DSGE MODEL

A classic question in economics is: Why do prices take so long to respond to a monetary disturbance and why do real variables react so strongly? Mankiw, writing in the year 2000, maintained that an empirically successful explanation of monetary non-neutrality has confounded economists at least since David Hume wrote "Of Money" in 1752. Moreover, at the time that Mankiw was writing, it looked as though the question would remain unanswered for a considerable time to come. A reason that monetary DSGE models have been so successful in the past decade is that, with a combination of modest price and wage stickiness and various "real frictions," they roughly reproduce the evidence of monetary non-neutrality that had seemed so hard to match. The purpose of this section and the next two sections is to spell out the basis for this observation in detail. Inevitably, doing this requires a model that is more complicated than the various versions of the simple model studied in the previous sections. In describing the model in this section, we explain the rationale for each departure from the simple model.

The model developed here is a version of the one in CEE. We describe the objectives and constraints of the agents in the model, and leave the derivation of the equilibrium conditions to the technical appendix. This model includes monetary policy shocks, so that it can be used to address the monetary non-neutrality question. In addition, the model includes two technology shocks. As a further check on the model, we follow ACEL in also evaluating the model's ability to match the estimated dynamic response of economic variables to the two technology shocks.

4.1 Goods production

An aggregate homogeneous good is produced using the technology, (Eq. 5). The first-order condition of the representative, competitive producer of the homogeneous good is given by Eq. (6). Substituting this first-order condition back into Eq. (5) yields the restriction across prices, Eq. (7). Each intermediate good, $i \in (0,1)$, is produced by a monopolist who treats Eq. (6) as its demand curve. The intermediate good producer takes the aggregate quantities, P_t and Y_t, as exogenous.

We use a production function for intermediate good producers that is standard in the literature. It does not use materials inputs, but it does use the services of capital, $K_{i,t}$:

$$Y_{i,t} = (z_t H_{i,t})^{1-\alpha} K_{i,t}^{\alpha} - z_t^{+}\varphi. \tag{59}$$

Here, z_t is a technology shock whose logarithmic first difference has a positive mean and φ denotes a fixed production cost. The economy has two sources of growth: the positive drift in log (z_t) and a positive drift in log (Ψ_t), where Ψ_t is the state of an investment specific technology shock discussed later. The object, z_t^{+}, in Eq. (59) is defined as follows:

$$z_t^{+} = \Psi_t^{\frac{\alpha}{1-\alpha}} z_t.$$

Along a nonstochastic steady-state growth path, Y_t/z_t^{+} and $Y_{i,t}/z_t^{+}$ converge to constants.

The two shocks, z_t and Ψ_t, are specified to be unit root processes in order to be consistent with the assumptions we use in our VAR analysis to identify the dynamic response of the economy to neutral and investment specific technology shocks. We adopt the following time series representations for the shocks:

$$\Delta \log z_t = \mu_z + \varepsilon_t^z, E\left(\varepsilon_t^z\right)^2 = \sigma_z^2 \tag{60}$$

$$\Delta \log \Psi_t = (1-\rho_\Psi)\mu_\Psi + \rho_\Psi \Delta \log \Psi_{t-1} + \varepsilon_t^\Psi, E\left(\varepsilon_t^\Psi\right)^2 = \sigma_\Psi^2. \tag{61}$$

Our assumption that the neutral technology shock follows a random walk with drift closely matches the finding in Smets and Wouters (2007) who estimated log z_t to be highly autocorrelated. The empirical analysis of Prescott (1986) also supported the notion that log z_t is a random walk with drift. Finally, Fernald (2009) constructed a direct estimate of total factor productivity growth for the business sector. The first-order autocorrelation of quarterly observations covering the period 1947Q2 to 2009Q3 is 0.0034, consistent with the idea of a random walk.

We assume that there is no entry or exit by intermediate good producers. The no entry assumption would be implausible if firms enjoyed large and persistent profits. The fixed cost in Eq. (59) is introduced to minimize the incentive to enter. We set φ so that intermediate good producer profits are zero in steady state. This requires that the fixed cost grow at the same rate as the growth rate of economic output, and this is why φ is multiplied by z_t^+ in Eq. (59). A potential empirical advantage of including fixed costs of production is that, by introducing some increasing returns to scale, the model can in principle account for evidence that labor productivity rises in the wake of a positive monetary policy shock.

In Eq. (59), $H_{i,t}$ denotes homogeneous labor services hired by the i^{th} intermediate good producer. Firms must borrow the wage bill. We follow CEE in supposing that firms borrow the entire wage bill (i.e., $\psi = 1$ and $\nu_t = 0$ in Eq. 9) so that the cost of one unit of labor is given by

$$W_t R_t. \tag{62}$$

Here, W_t denotes the aggregate wage rate and R_t denotes the gross nominal interest rate on working capital loans. The assumption that firms require working capital was introduced by CEE as a way to help dampen the rise in inflation after an expansionary shock to monetary policy. An expansionary shock to monetary policy drives R_t down and, with other things the same, this reduces firm marginal cost. Inflation is dampened because marginal cost is the key input into firms' price-setting decisions. Indirect evidence consistent with the working capital assumption includes the frequently found VAR-based results, suggesting that inflation drops for a little while after a positive monetary policy shock. It is hard to think of an alternative to the working capital

assumption to explain this evidence, apart from the possibility that the estimated response reflects some kind of econometric specification error.[40]

Another motivation for treating interest rates as part of the cost of production has to do with Ball's (1994) "dis-inflationary boom" critique models that do not include interest rates in costs. Ball's critique focuses on the Phillips curve in Eq. (30), which we reproduce here for convenience:

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa_p \hat{s}_t,$$

where $\hat{\pi}_t$ and $\hat{s}_t$ denote inflation and marginal cost, respectively. Also, $\kappa_p > 0$ is a reduced form parameter and β is slightly less than unity. According to the Phillips curve, if the monetary authority announces it will fight inflation by strategies that (plausibly) bring down future inflation more than present inflation, then $\hat{s}_t$ must jump. In simple models $\hat{s}_t$ is directly related to the volume of output (see, e.g., Eq. 34). High output requires more intense utilization of scarce resources, their price goes up, driving up marginal cost, $\hat{s}_t$. Ball (1994) criticized theories that do not include the interest rate in marginal cost on the grounds that we do not observe booms at the start of disinflations. Including the interest rate in marginal cost potentially avoids the Ball critique because the high $\hat{s}_t$ may simply reflect the high interest rate that corresponds to the disinflationary policy, and not higher output.

We adopt the Calvo model of price frictions. With probability ξ_p, the intermediate good firm cannot reoptimize its price, in which case it is assumed to set its price according to the following rule:[41]

$$P_{i,t} = \pi P_{i,t-1}. \tag{63}$$

Note that in steady state, firms that do not optimize their prices raise prices at the general rate of inflation. Firms that optimize their prices in a steady-state growth path raise their prices by the same amount. This is why there is no price dispersion in steady state. According to the discussion near Eq. (29), the fact that we analyze the first-order approximation of a DSGE model in a neighborhood of steady state means that we can impose the analog of $p_t^* = 1$.

With probability $1 - \xi_p$ the intermediate good firm can reoptimize its price. Apart from the fixed cost, the i^{th} intermediate good producer's profits are the analog of Eq. (13):

$$E_t \sum_{j=0}^{\infty} \beta^j \upsilon_{t+j} \left[P_{i,t+j} Y_{i,t+j} - s_{t+j} P_{t+j} Y_{i,t+j} \right],$$

[40] This possibility was suggested by Sims (1992) and explored further in Christiano et al. (1999). See also Bernanke, Boivin, and Eliasz (2005).

[41] Equation (63) excludes the possibility that firms index to past inflation. We discuss the reason for this specification in Section 6.2.2.

where s_t denotes the marginal cost of production, denominated in units of the homogeneous good. The object, s_t, is a function only of the costs of capital and labor, and is described in Section C of the technical appendix . Marginal cost is independent of the level of $Y_{i,t}$ because of the linear homogeneity of the first expression on the right of Eq. (59). The first-order necessary conditions associated with this optimization problem are reported in section E of the technical appendix.

Goods market clearing dictates that the homogeneous output good is allocated among alternative uses as follows:

$$Y_t = G_t + C_t + \widetilde{I}_t. \tag{64}$$

Here, C_t denotes household consumption, G_t denotes exogenous government consumption, and $\tilde{I}_t$ is a homogenous investment good which is defined as follows:

$$\widetilde{I}_t = \frac{1}{\Psi_t}(I_t + a(u_t)\bar{K}_t). \tag{65}$$

The investment goods, I_t, are used by households to add to the physical stock of capital, $\bar{K}_t$.[42] The remaining investment goods are used to cover maintenance costs, $a(u_t)\ \bar{K}_t$, arising from capital utilization, u_t. The cost function, $a\ (\cdot)$, is increasing and convex, and has the property that in steady state, $u_t = 1$ and $a(1) = 0$. The relationship between the utilization of capital, u_t, capital services, K_t, and the physical stock of capital, $\bar{K}_t$, is as follows:

$$K_t = u_t\bar{K}_t.$$

The investment and capital utilization decisions are discussed in Section 4.2. See Section 4.4 for the functional form of the capital utilization cost function. Finally, Ψ_t in Eq. (65) denotes the unit root investment specific technology shock defined in Eq. (61).

4.2 Households

In the model, households supply the factors of production, labor and capital. The model incorporates Calvo-style wage setting frictions along the lines spelled out in Erceg, Henderson, and Levin (2000). Because wages are an important component of costs, wage-setting frictions help slow the response of inflation to a monetary policy shock. As in the case of prices, wage-setting frictions require that there be market power. To ensure that this market power is suffused through the economy and not, for example, concentrated in the hands of a single labor union, we adopt the framework that is now standard in monetary DSGE models. In particular, we adopt a variant

[42] The notation, I_t, used here should not be confused with materials inputs in Section 2. Our medium-sized DSGE model does not include materials inputs.

of the model in Erceg et al. (2000) by using the analog of the Dixit-Stiglitz type framework used to model price-setting frictions. The assumption that prices are set by producers of specialized goods appears here in the form of the assumption that there are many different specialized labor inputs, $h_{j,t}$, for $j \in (0,1)$. There is a single monopolist that sets the wage for each type, j, of labor service. However, that monopolist's market power is severely limited by the presence of other labor services, $j' \neq j$, that are substitutable for $h_{j,t}$.

The variant of the Erceg et al. (2000) model that we work with follows the discussion in Section 2.3 in supposing that labor is indivisible: people work either full time or not at all.[43] That is, $h_{j,t}$ represents a quantity of people and not, say, the number of hours worked by a representative worker.

Section 4.2.1 below discusses the interaction between households and the labor market. The Section 4.2.2 discusses the monopoly wage-setting problem in the model. Section 4.2.3 discusses the representative household's capital accumulation decision, and Section 4.2.4 states the representative household's optimization problem.

4.2.1 Households and the labor market

The "labor" hired by firms in the goods-producing sector is interpreted as a homogeneous factor of production, H_t, supplied by "labor contractors." Labor contractors produce H_t by combining a range of differentiated labor inputs, $h_{t,j}$, using the following linear homogeneous technology:

$$H_t = \left[\int_0^1 (h_{t,j})^{\frac{1}{\lambda_w}} dj\right]^{\lambda_w}, \lambda_w > 1.$$

Labor contractors are perfectly competitive and take the wage rate, W_t, of H_t as given. They also take the wage rate, $W_{t,j}$, of the j^{th} labor type as given. Contractors choose inputs and outputs to maximize profits,

$$W_t H_t - \int_0^1 W_{t,j} h_{t,j} dj.$$

The first order necessary condition for optimization is given by:

$$h_{t,j} = \left(\frac{W_t}{W_{t,j}}\right)^{\frac{\lambda_w}{\lambda_w - 1}} H_t. \tag{66}$$

Substituting the latter back into the labor aggregator function and rearranging, we obtain:

[43] Our approach follows the one in Gali (2010a).

$$W_t = \left[\int_0^1 W_{t,j}^{\frac{1}{1-\lambda_w}} dj\right]^{1-\lambda_w}. \tag{67}$$

Differentiated labor is supplied by a large number of identical households. The representative household has many members corresponding to each type, j, of labor. Each worker of type j has an index, l, distributed uniformly over the unit interval, [0,1], which indicates that worker's aversion to work. A type j worker with index l experiences utility:

$$\log\left(c_t^e - bC_{t-1}\right) - l^{\phi}, \phi > 0,$$

if employed and

$$\log\left(c_t^{ne} - bC_{t-1}\right),$$

if not employed. When $b > 0$ the worker's marginal utility of current consumption is an increasing function of the household's consumption in the previous period. Given the additive separability of consumption and employment in utility, the efficient allocation of consumption across workers within the household implies:[44]

$$c_t^e = c_t^{ne} = C_t.$$

The quantity of the j^{th} type of labor supplied by the representative household, $h_{t,j}$, is determined by Eq. (66). We suppose the household sends j-type workers with $0 \leq l \leq h_{t,j}$ to work and keeps those with $l > h_{t,j}$ out of the labor force. The equally weighted integral of utility over all $l \in [0,1]$ workers is

$$\log\left(C_t - bC_{t-1}\right) - A\frac{h_{t,j}^{1+\phi}}{1+\phi}.$$

Aggregate household utility also integrates over the unit measure of j-type workers:

$$\log\left(C_t - bC_{t-1}\right) - A\int_0^1 \frac{h_{t,j}^{1+\phi}}{1+\phi} dj. \tag{68}$$

Next we explain how $h_{t,j}$ is determined and how the household chooses C_t.

The wage rate of the j^{th} type of labor, $W_{t,j}$, is determined outside the representative household by a monopoly union that represents all j-type workers across all households. The union's problem is discussed in Section 4.2.2.

The presence of $b > 0$ in Eq. (68) is motivated by VAR-based evidence like that displayed in section 6 below, which suggests that an expansionary monetary policy shock triggers (i) a hump-shaped response in consumption and (ii) a persistent

[44] For an environment in which perfect insurance is not feasible, see CTW.

reduction in the real rate of interest.[45] With $b = 0$ and a utility function separable in labor and consumption like the one in Eq. (68), (i) and (ii) are difficult to reconcile. An expansionary monetary policy shock that triggers an increase in expected future consumption would be associated with a rise in the real rate of interest, not a fall. Alternatively, a fall in the real interest rate would cause people to rearrange consumption intertemporally, so that consumption is relatively high right after the monetary shock and low later. Intuitively, one can reconcile (i) and (ii) by supposing the marginal utility of consumption is inversely proportional not to the level of consumption, but to its derivative. To see this, it is useful to recall the familiar intertemporal Euler equation implied by household optimization (see, e.g., Eq. 4):

$$\beta E_t \frac{u_{c,t+1}}{u_{c,t}} \frac{R_t}{\pi_{t+1}} = 1.$$

Here, $u_{c,t}$ denotes the marginal utility of consumption at time t. From this expression, we see that a low R_t/π_{t+1} tends to produce a high $u_{c,t+1}/u_{c,t}$; that is, a rising trajectory for the marginal utility of consumption. This illustrates the problematic implication of the model when $u_{c,t}$ is inversely proportional to C_t as in Eq. (68) with $b = 0$. To fix this implication we need a model change with the property that a rising $u_{c,t}$ path implies hump-shaped consumption. A hump-shaped consumption path corresponds to a scenario in which the slope of the consumption path is falling, suggesting that (i) and (ii) can be reconciled if $u_{c,t}$ is proportional to the slope of consumption. The notion that marginal utility is inversely proportional to the slope of consumption corresponds loosely to $b > 0$.[46] The fact that (i) and (ii) can be reconciled with the assumption of habit persistence is of special interest, because there is evidence from other sources that also favors the assumption of habit persistence; for example, in asset pricing (see, e.g., Boldrin, Christiano, & Fisher, 2001; Constantinides, 1990) and growth (see Carroll, Overland, & Weil, 1997, 2000). In addition, there may be a solid foundation in psychology for this specification of preferences.[47]

[45] The earliest published statement of the idea that $b > 0$ can help account for (i) and (ii) that we are aware of is Fuhrer (2000).

[46] In particular, suppose first that lagged consumption in Eq. (68) represents aggregate, economy-wide consumption and $b > 0$. This corresponds to the so-called "external habit" case, where it is the lagged consumption of others that enters utility. In that case, the marginal utility of household C_t is $1/(C_t - bC_{t-1})$, which corresponds to the inverse of the slope of the consumption path, at least if b is large enough. In our model we think of C_{t-1} as corresponding to the household's own lagged consumption (that is why we use the same notation for current and lagged consumption), the so-called "internal habit" case. In this case, the marginal utility of C_t also involves future terms, in addition to the inverse of the slope of consumption from $t = 1$ to t. The intuition described in the text, which implicitly assumed external habit, also applies roughly to the internal habit case that we consider.

[47] Anyone who has gone swimming has had the experience of habit persistence. It is usually very hard at first to jump into a swimming pool because it seems so cold. The swimmer who jumps (or is pushed) into the water after much procrastination initially experiences a tremendous shock with the sudden drop in temperature. However, after only a few minutes the new, lower temperature is perfectly comfortable. In this way, the lagged temperature seems to influence one's experience of current temperature, as in habit persistence.

The logic associated with the previous intertemporal Euler equation suggests that there are other approaches that can at least go part way in reconciling (i) and (ii). For example, Guerron-Quintana (2008) showed that nonseparability between consumption and labor in Eq. (68) can help reconcile (i) and (ii). He pointed out that if the marginal utility of consumption is an increasing function of labor and the model predicts that employment rises in a hump shape after an expansionary monetary shock, then it is possible that consumption rises in a hump shape.

4.2.2 Wages, employment and monopoly unions

We turn now to a discussion of the monopoly union that sets the wage of j-type workers. In each period, the monopoly union must satisfy its demand curve, (Eq. 66), and it faces Calvo frictions in the setting of $W_{t,j}$. With probability $1-\xi_w$ the union can optimize the wage and with the complementary probability, ξ_w, it cannot. In the latter case, we suppose that the nominal wage rate is set as follows:

$$W_{j,t+1} = \widetilde{\pi}_{w,t+1} W_{j,t} \tag{69}$$

$$\widetilde{\pi}_{w,t+1} = \pi_t^{\kappa_w} \pi^{(1-\kappa_w)} \mu_{z^+}, \tag{70}$$

where $\kappa_w \in (0, 1)$. With this specification, the wage of each type j labor is the same in the steady state. Because the union problem has no state variable, all unions with the opportunity to reoptimize in the current period face the same problem. In particular, such a union chooses the current value of the wage, $\widetilde{W}_t$, to maximize:

$$E_t \sum_{i=0}^{\infty} (\beta\xi_w)^i \left[\upsilon_{t+i} \widetilde{W}^t_{t+i} h^t_{t+i} - A_L \frac{\left(h^t_{t+i}\right)^{1+\phi}}{(1+\phi)} \right]. \tag{71}$$

Here, h^t_{t+i} and $\widetilde{W}^t_{t+i}$ denote the quantity of workers employed and their wage rate, in period $t+i$, of a union that has an opportunity to reoptimize the wage in period t and does not reoptimize again in periods $t+1, \ldots, t+i$. Also, υ_{t+i} denotes the marginal value assigned by the representative household to the wage.[48] The union treats υ_t as an exogenous variable. In the expression (71), ξ_w appears in the discounting because the union's period t decision only impacts on future histories in which it cannot reoptimize its wage.

Optimization by all labor unions leads to a simple equilibrium condition, when the variables are linearized about the nonstochastic steady state.[49] The condition is:

[48] The object, υ_t, is the multiplier on the household budget constraint in the Lagrangian representation of the problem.
[49] The details of the derivation are explained in Section G of the technical appendix.

$$\Delta_{\kappa_w}\hat{\pi}_{w,t} = \frac{\kappa}{1+\phi\dfrac{\lambda_w}{\lambda_w-1}}\left(\overbrace{-\hat{\psi}_{z^+,t}+\phi\hat{H}_t}^{\text{scaled labor cost of marginal worker}} - \overbrace{\hat{\hat{w}}_t}^{\text{scaled real wage}}\right) + \beta\Delta_{\kappa_w}\hat{\pi}_{w,t+1}, \tag{72}$$

where

$$\kappa = \frac{(1-\xi_w)(1-\beta\xi_w)}{\xi_w}.$$

In Eq. (72), $\hat{\pi}_{w,t}$ is the gross growth rate in the nominal wage rate, expressed in percent deviation from steady state. Also, $\hat{\psi}_{z^+,t}$ represents the percent deviation of the scaled multiplier, $\psi_{z^+,t}$, from its steady-state value. The scaled multiplier is defined as follows:

$$\psi_{z^+,t} \equiv \upsilon_t P_t z_t^+,$$

where υ_t is the multiplier on the household budget constraint. The first two terms inside the parentheses in Eq. (72) correspond to the marginal cost of labor and the third term, $\hat{\hat{w}}_t$, corresponds to the real wage. Both the marginal cost of labor and the real wage have been scaled by z_t^+. Expression (72) has a simple interpretation. The first term in parentheses is related to the cost of working by the marginal worker. When this (scaled) cost exceeds the (scaled) real wage, $\hat{\hat{w}}_t$, then the monopoly unions currently setting wages place upward pressure on wage inflation. The coefficient multiplying the term in parentheses is also interesting. If the degree of wage and price stickiness are the same; that is, $\xi_w = \xi_p$, then κ takes on the same value as κ_p, the analog of κ in the price Phillips curve, Eq. (35). In this case, the slope of the price Phillips curve in terms of marginal cost is bigger than the slope of the wage Phillips curve, Eq. (72). This reflects that in the slope of the wage Phillips curve, κ is divided by:

$$1+\phi\frac{\lambda_w}{\lambda_w-1} > 1.$$

According to this expression, the slope of the wage Phillips curve is smaller if the elasticity of demand for labor, $\lambda_w/(\lambda_w - 1)$ is large and/or if the marginal cost of work, *MRS*, is sharply increasing in work (i.e., ϕ is large). The intuition for this is as follows. Suppose the j^{th} monopoly union contemplates a particular rise in the nominal wage, for whatever reason. Consider a given slope of the demand for labor. The rise in the wage implies a lower quantity of labor demanded. The steeper the marginal cost curve is, the greater the implied drop in marginal cost. Now consider a given slope of marginal cost. The flatter the slope of demand for the j^{th} type of labor is, the larger the drop in the quantity of labor demanded in response to the given contemplated rise

in the wage. Given the upward-sloping marginal cost curve, this also implies a large fall in marginal cost. Thus, the monopoly union that contemplates a given rise in the wage rate anticipates a larger drop in marginal cost to the extent that the demand curve is elastic and/or the marginal cost curve is steep. But, with other things the same, low marginal cost reduces the incentive for a monopolist to raise its price (i.e., the wage in this case). These considerations are absent in our price Phillips curve, Eq. (35), because marginal cost is constant (i.e., the analog of ϕ is zero).[50]

4.2.3 Capital accumulation

The household owns the economy's physical stock of capital, sets the utilization rate of capital, and rents out the services of capital in a competitive market. The household accumulates capital using the following technology:

$$\bar{K}_{t+1} = (1-\delta)\bar{K}_t + F(I_t, I_{t-1}) + \Delta_t, \tag{73}$$

where Δ_t denotes physical capital purchased in a competitive market from other households. Since all households are the same in terms of capital accumulation decisions, $\Delta_t = 0$ in equilibrium. We nevertheless include Δ_t so that we can assign a price to installed capital. In Eq. (73), $\delta \in [0, 1]$ and we use the specification suggested in CEE:

$$F(I_t, I_{t-1}) = \left(1 - S\left(\frac{I_t}{I_{t-1}}\right)\right) I_t, \tag{74}$$

where the functional form, S, that we use is described in Section 4.4. In Eq. (74), $S = S' = 0$ and $S'' > 0$ along a nonstochastic steady state growth path. Here, S' and S'' denote the first and second derivatives, respectively, of S.

Let $P_t P_{k',t}$ denote the nominal market price of Δ_t. For each unit of $\bar{K}_{t+1}$ acquired in period t, the household receives X^k_{t+1} in net cash payments in period $t+1$:

$$X^k_{t+1} = u_{t+1} P_{t+1} r^k_{t+1} - \frac{P_{t+1}}{\Psi_{t+1}} a(u_{t+1}). \tag{75}$$

The first term is the gross nominal period $t+1$ rental income from a unit of $\bar{K}_{t+1}$. The second term represents the cost of capital utilization, $a(u_{t+1})P_{t+1}/\Psi_{t+1}$. Here, P_{t+1}/Ψ_{t+1} is the nominal price of the investment goods absorbed by capital utilization. That P_{t+1}/Ψ_{t+1} is the equilibrium market price of investment goods follows from the technology specified in Eqs. (64) and (65), and the assumption that investment goods are produced from homogeneous output goods by competitive firms.

The introduction of variable capital utilization is motivated by a desire to explain the slow response of inflation to a monetary policy shock. In any model prices are

[50] This intuition for why the slope of the wage Phillips curve is flatter with elastic labor demand and/or steep marginal cost is the same as the intuition that firm-specific capital flattens the price Phillips curve (see, e.g., ACEL; Christiano, 2004; de Walque, Smets, & Wouters, 2006; Sveen & Weinke, 2005; Woodford, 2004).

heavily influenced by costs. Costs in turn are influenced by the elasticity of the factors of production. If factors can be rapidly expanded with a small rise in cost, then inflation will not rise much after a monetary policy shock. Allowing for variable capital utilization is a way to make the services of capital elastic. If there is very little curvature in the *a* function, then households are able to expand capital services without much increase in cost.

The form of the investment adjustment costs in Eq. (73) is motivated by a desire to reproduce VAR-based evidence that investment has a hump-shaped response to a monetary policy shock. Alternative specifications include $F \equiv I_t$ and

$$F = I_t - \frac{S''}{2}\left(\frac{I_t}{\bar{K}_t} - \delta\right)^2 \bar{K}_t. \tag{76}$$

Specification (76) has a long history in macroeconomics, and has been in use since at least Lucas and Prescott (1971). To understand why DSGE models generally use the adjustment cost specification in Eq. (74) rather than Eq. (76), it is useful to define the rate of return on investment:

$$R^k_{t+1} = \frac{x^k_{t+1} + \left[1 - \delta + S''\left(\frac{I_{t+1}}{\bar{K}_{t+1}} - \delta\right)\frac{I_{t+1}}{\bar{K}_{t+1}} - \frac{S''}{2}\left(\frac{I_{t+1}}{\bar{K}_{t+1}} - \delta\right)^2\right] P_{k',t+1}}{P_{k',t}}. \tag{77}$$

The numerator is the one-period payoff from an extra unit of $\bar{K}_{t+1}$, and the denominator is the corresponding cost, both in consumption units. In Eq. (77), $x^k_{t+1} \equiv X^k_{t+1}/P_{t+1}$ denotes the earnings net of costs. The term in square brackets is the quantity of additional $\bar{K}_{t+2}$ made possible by the additional unit of $\bar{K}_{t+1}$. This is composed of the undepreciated part of $\bar{K}_{t+1}$ left over after production in period $t+1$, plus the impact of $\bar{K}_{t+1}$ on $\bar{K}_{t+2}$ via the adjustment costs. The object in square brackets is converted to consumption units using $P_{k',t+1}$, which is the market price of $\bar{K}_{t+2}$ denominated in consumption goods. Finally, the denominator is the price of the extra unit of $\bar{K}_{t+1}$.

The price of extra capital in competitive markets corresponds to the marginal cost of production. Thus,

$$\begin{aligned} P_{k',t} &= -\frac{dC_t}{d\bar{K}_{t+1}} = -\frac{dC_t}{dI_t} \times \frac{dI_t}{d\bar{K}_{t+1}} \\ &= \frac{1}{\frac{d\bar{K}_{t+1}}{dI_t}} = \begin{cases} 1 & \text{When } F \text{ is } I \\ \dfrac{1}{1 - S'' \times \left(\dfrac{I_t}{\bar{K}_t} - \delta\right)} & \text{When } F \text{ is as in (76)}, \end{cases} \end{aligned} \tag{78}$$

where we ignore Ψ_t for now ($\Psi_t \equiv 1$). The derivatives in the first line correspond to marginal rates of technical transformation. The marginal rate of technical transformation between consumption and investment is implicit in Eqs. (64) and (65). The marginal rate of technical transformation between I_t and $\bar{K}_{t+1}$ is given by the capital accumulation equation. The relation in the second line of Eq. (78) is referred to as "Tobin's q" relation, where Tobin's q here corresponds to $P_{k',t}$. This is the market value of capital (i.e., the marginal cost of capital under our assumption that markets are competitive) divided by the price of investment goods. Here, q can differ from unity due to the investment adjustment costs.

We are now in a position to convey the intuition about why DSGE models have generally abandoned the specification in Eq. (76) in favor of Eq. (73). The key reason has to do with VAR-based evidence (see Section 6 below) that suggests the real interest rate falls persistently after a positive monetary policy shock, while investment responds in a hump-shaped pattern. Any model that is capable of producing this type of response will have the property that the real return on capital, Eq. (77) — for arbitrage reasons — also falls after an expansionary monetary policy shock. Suppose, to begin, that $S'' = 0$, so that there are no adjustment costs at all and $P_{k',t} = 1$. In this case, the only component in R_t^k that can fall is x_{t+1}^k, which is dominated by the marginal product of capital. That is, approximately, the rate of return on capital is

$$(1-\alpha)K_{t+1}^{\alpha-1}H_{t+1}^{1-\alpha} + 1 - \delta.$$

In steady state this object equals $1/\beta$ (ignoring growth), which is roughly 1.03 in annual terms. At the same time, the object, $1-\delta$, is roughly 0.9 in annual terms, so that the endogenous part of the rate of return of capital is a very small part of that rate of return. As a result, any given drop in the return on capital requires a very large percentage drop in the endogenous part, $K_{t+1}^{\alpha-1}H_{t+1}^{1-\alpha}$. An expansion in investment can bring about this drop, but it has to be a very substantial surge. To see why investment must expand so much, note first that the endogenous part of the rate of return is not only small, but the capital stock receives a weight substantially less than unity in that expression. Second, a model that successfully reproduces the VAR-based evidence that employment rises after a positive monetary policy shock implies that hours worked rises. This pushes the endogenous component of the rate of return up, increasing the burden on the capital stock to bring the rate of return on investment down. For these reasons, models without adjustment costs generally imply a counterfactually strong surge in investment in the wake of a positive shock to monetary policy.

With $S'' > 0$ the endogenous component of the rate of return on capital is much larger. However, in practice models that adopt the adjustment cost specification, Eq. (76), generally imply that the biggest investment response occurs in the period of the shock, and not later. To gain intuition into why this is so, suppose the contrary:

that investment does exhibit a hump-shape response in investment. Equation (78) implies a similar hump-shape pattern in the price of capital, $P_{k',t}$.[51] This is because $P_{k',t}$ is primarily determined by the contemporaneous flow of investment. So, under our supposition about the investment response, a positive monetary policy shock generates a rise in $P_{k',t+1}/P_{k',t}$ over at least several periods in the future. According to Eq. (77), the anticipated future capital gains create an incentive to invest right away. Thus, households would be induced to substitute away from a hump-shaped response, toward one in which the immediate response is much stronger. In practice, this means that in equilibrium, the biggest response of investment occurs in the period of the shock, with later responses converging to zero.

The adjustment costs in Eq. (74) do have the implication that investment responds in a hump-shaped manner. According to Eq. (74), a quick rise in investment over its previous level is costly.

There are other reasons to take the specification in Eq. (74) seriously. Lucca (2006) and Matsuyama (1984) have described interesting theoretical foundations that produce Eq. (74) as a reduced form. For example, in Matsuyama shifting production between consumption and capital goods involves a learning-by-doing process, which makes quick movements in either direction expensive. Also, Matsuyama explains why the abundance of empirical evidence that appears to reject Eq. (76) may be consistent with Eq. (74). Consistent with Eq. (74), Topel and Rosen (1988) argue that data on housing construction cannot be understood without using a cost function that involves the change in the flow of housing construction.

4.2.4 Household optimization problem

The household's period t budget constraint is as follows:

$$P_t\left(C_t+\frac{1}{\Psi_t}I_t\right)+B_{t+1}+P_tP_{k',t}\Delta_t\leq\int_0^1 W_{t,j}h_{t,j}dj+X_t^k\bar{K}_t+R_{t-1}B_t, \qquad (79)$$

where $W_{t,j}$ represents the wage earned by the household, B_{t+1} denotes the quantity of risk-free bonds purchased by the household, and R_t denotes the gross nominal interest rate on bonds purchased in period $t-1$, which pay off in period t. The household's problem is to select sequences, $\{C_t, I_t, \Delta_t, B_{t+1}, \bar{K}_{t+1}\}$, to maximize Eq. (68) subject to the wage process selected by the monopoly unions, Eqs. (73), (75), and (79).

[51] Note from Eq. (78) that the price of capital increases as investment rises above its level in steady state, which is the level required to just meet the depreciation in the capital stock. Our assertion that the price of capital follows the same hump-shaped pattern as investment after a positive monetary policy shock reflects our implicit assumption that the shock occurs when the economy is in a steady state. This will be true on average, but not at each date.

4.3 Fiscal and monetary authorities and equilibrium

We suppose that monetary policy follows a Taylor rule of the following form:

$$\log\left(\frac{R_t}{R}\right) = \rho_R \log\left(\frac{R_{t-1}}{R}\right) + (1-\rho_R)\left[r_\pi \log\left(\frac{\pi_{t+1}}{\pi}\right) + r_y \log\left(\frac{gdp_t}{gdp}\right)\right] + \varepsilon_{R,t}, \quad (80)$$

where $\varepsilon_{R,t}$ denotes an iid shock to monetary policy. As in CEE and ACEL, we assume that the period t realization of $\varepsilon_{R,t}$ is not included in the period t information set of the agents in our model. This ensures that our model satisfies the restrictions used in the VAR analysis to identify a monetary policy shock. In Eq. (80), gdp_t denotes scaled real GDP defined as follows:

$$gdp_t = \frac{G_t + C_t + I_t/\Psi_t}{z_t^+}. \quad (81)$$

We adopt the model of government consumption suggested in Christiano and Eichenbaum (1992a):

$$G_t = g z_t^+.$$

In principle, g could be a random variable, although our focus in this paper is just on monetary policy and technology shocks. So, we set g to a constant. Lump-sum transfers are assumed to balance the government budget.

An equilibrium is a stochastic process for the prices and quantities with the property that the household and firm problems are satisfied, and goods and labor markets clear.

4.4 Adjustment cost functions

We adopt the following functional forms. The capacity utilization cost function is

$$a(u) = 0.5b\sigma_a u^2 + b(1-\sigma_a)u + b((\sigma_a/2) - 1), \quad (82)$$

where b is selected so that $a(1) = a'(1) = 0$ in steady state and σ_a is a parameter that controls the curvature of the cost function. The closer σ_a is to zero, the less curvature there is and the easier it is to change utilization. The investment adjustment cost function takes the following form:

$$\begin{aligned} S(x_t) &= \frac{1}{2}\left\{\exp\left[\sqrt{S''}(x_t - \mu_{z^+}\mu_\Psi)\right] + \exp\left[-\sqrt{S''}(x_t - \mu_{z^+}\mu_\Psi)\right] - 2\right\}, \\ &= 0, x = \mu_{z^+}\mu_\Psi. \end{aligned} \quad (83)$$

where $x_t = I_t/I_{t-1}$ and $\mu_{z^+}\mu_\Psi$ is the growth rate of investment in steady state. With this adjustment cost function, $S(\mu_{z^+}\mu_\Psi) = S'(\mu_{z^+}\mu_\Psi) = 0$. Also, $S'' > 0$ is a parameter having the property that it is the second derivative of $S(x_t)$ evaluated at $x_t = \mu_{z^+}\mu_\psi$. Because of the nature of the above adjustment cost functions, the curvature parameters have no impact on the model's steady state.

5. ESTIMATION STRATEGY

Our estimation strategy is a Bayesian version of the two-step impulse response matching approach applied by Rotemberg and Woodford (1997) and CEE. We begin with a discussion of the two steps. After that, we discuss the computation of a particular weighting matrix used in the analysis.

5.1 VAR step

We estimate the dynamic responses of a set of aggregate variables to three shocks, using standard VAR methods. The three shocks are the monetary policy shock; the innovation to the permanent technology shock, z_t; and the innovation to the investment specific technology shock, Ψ_t. The estimated contemporaneous and 14 lagged responses in each of $N = 9$ macroeconomic variables to the three shocks are stacked in a vector, $\hat{\psi}$. These macroeconomic variables are a subset of the variables that appear in the VAR. The additional variables in our VAR pertain to the labor market. We use this augmented VAR to facilitate comparison between the analysis in this chapter and in other research where we integrate labor market frictions into a monetary DSGE model.[52] We denote the vector of variables in the VAR by Y_t, where[53]

$$\underbrace{Y_t}_{14\times 1} = \begin{pmatrix} \Delta \ln(\text{relative price of investment}_t) \\ \Delta \ln(\text{real } GDP_t/\text{hours}_t) \\ \Delta \ln(GDP \text{ deflator}_t) \\ \text{unemployment rate}_t \\ \text{capacity utilization}_t \\ \ln(\text{hours}_t) \\ \ln(\text{real } GDP_t/\text{hours}_t) - \ln(W_t/P_t) \\ \ln(\text{nominal } C_t/\text{nominal } GDP_t) \\ \ln(\text{nominal } I_t/\text{nominal } GDP_t) \\ \text{vacancies}_t \\ \text{job separation rate}_t \\ \text{job finding rate}_t \\ \log(\text{hours}_t/\text{labor force}_t) \\ \text{Federal Funds Rate}_t \end{pmatrix}. \tag{84}$$

An extensive general review of identification in VARs appears in Christiano et al. (1999). The specific technical details of how we compute impulse response functions

[52] See Christiano, Trabandt and Walentin (2010b). The variables, GDP, hours, C, I and the labor force, are expressed in per capita terms.

[53] See Section A of the technical appendix for details about the data.

imposing the shock identification are reported in ACEL.[54] We estimate a two-lag VAR using quarterly data that are seasonally adjusted and cover the period 1951Q1 to 2008Q4. Our identification assumptions are as follows. The only variable that the monetary policy shock affects contemporaneously is the federal funds rate. We make two assumptions to identify the dynamic response to the technology shocks: (i) the only shocks that affect labor productivity in the long run are the two technology shocks, and (ii) the only shock that affects the price of investment relative to consumption is the innovation to the investment specific shock. All of these identification assumptions are satisfied in our model.

Our data set extends over a long range, while we estimate a single set of impulse response functions and model parameters. In effect, we suppose that there has been no parameter break over this long period. Whether or not there has been a break is a question that has been debated. For example, it has been argued that the parameters of the monetary policy rule have not been constant over this period. We do not review this debate here. Implicitly, our analysis sides with the conclusions of those that argue that the evidence of parameter breaks is not strong. For example, Sims and Zha (2006) argued that the evidence is consistent with the idea that monetary policy rule parameters have been unchanged over the sample. Christiano et al. (1999) argued that the evidence is consistent with the proposition that the dynamic effects of a monetary policy shock have not changed during this sample. Standard lag-length selection criteria led us to work with a VAR with 2 lags.[55]

The number of elements in $\hat{\psi}$ corresponds to the number of impulses estimated. Since we consider the contemporaneous and 14 lag responses in the impulses, there are in principle 3 (i.e., the number of shocks) times 9 (number of variables) times 15 (number of responses) = 405 elements in $\hat{\psi}$. However, we do not include in $\hat{\psi}$ the 8 contemporaneous responses to the monetary policy shock that are required to be zero by our monetary policy identifying assumption. Taking this into account, the vector $\hat{\psi}$ has 397 elements.

According to standard classical asymptotic sampling theory, when the number of observations, T, is large, we have

$$\sqrt{T}\left(\hat{\psi} - \psi(\theta_0)\right) \overset{a}{\sim} N(0, W(\theta_0, \zeta_0)),$$

[54] The identification assumption for the monetary policy shock by itself imposes no restriction on the VAR parameters. Similarly, Fisher (2006) showed that the identification assumptions for the technology shocks, when applied without simultaneously applying the monetary shock identification, also imposes no restriction on the VAR parameters. However, ACEL showed that when all the identification assumptions are imposed at the same time, then there are restrictions on the VAR parameters. We found that the test of the overidentifying restrictions on the VAR fails to reject the null hypothesis that the restrictions are satisfied at the 5% critical level.

[55] We considered VAR specifications with lag length 1, 2, ..., 12. The Schwartz and Hannan-Quinn criteria indicate that a single lag in the VAR is sufficient. The Akaike criterion indicates 12 lags, but we discounted that result.

where θ_0 represents the true values of the parameters that we estimate. The vector, ζ_0, denotes the true values of the parameters of the shocks that are in the model, but that we do not formally include in the analysis. We find it convenient to express the asymptotic distribution of $\hat{\psi}$ in the following form:

$$\hat{\psi} \overset{a}{\sim} N(\psi(\theta_0), V(\theta_0, \zeta_0, T)), \tag{85}$$

where

$$V(\theta_0, \zeta_0, T) \equiv \frac{W(\theta_0, \zeta_0)}{T}.$$

5.2 Impulse response matching step

In the second step of our analysis, we treat $\hat{\psi}$ as "data" and we choose a value of θ to make $\psi(\theta)$ as close as possible to $\hat{\psi}$. As discussed in Section 3.3.3 and following Kim (2002), we refer to our strategy as a limited information Bayesian approach. This interpretation uses Eq. (85) to define an approximate likelihood of the data, $\hat{\psi}$, as a function of θ :

$$\begin{aligned} f\left(\hat{\psi}|\theta\right) &= \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} |V(\theta_0, \zeta_0, T)|^{-\frac{1}{2}} \\ \times \exp&\left[-\frac{1}{2}\left(\hat{\psi} - \psi(\theta)\right)' V(\theta_0, \zeta_0, T)^{-1}\left(\hat{\psi} - \psi(\theta)\right)\right]. \end{aligned} \tag{86}$$

In Eq. (86), N denotes the number of elements in $\hat{\psi}$. As we explain next, we treat the value of $V(\theta_0,\zeta_0,T)$ as a known object. Under these circumstances, the value of θ that maximizes the above function represents an approximate maximum likelihood estimator of θ. It is approximate for two reasons: (i) the central limit theorem underlying Eq. (85) only holds exactly as $T \to \infty$ and (ii) the value of $V(\theta_0, \zeta_0, T)$ that we use is guaranteed to be correct only for $T \to \infty$.

Treating the function, f, as the likelihood of $\hat{\psi}$, it follows that the Bayesian posterior of θ conditional on $\hat{\psi}$ and $V(\theta_0, \zeta_0, T)$ is

$$f\left(\theta|\hat{\psi}\right) = \frac{f\left(\hat{\psi}|\theta\right)p(\theta)}{f\left(\hat{\psi}\right)}, \tag{87}$$

where $p(\theta)$ denotes the priors on θ and $f(\hat{\psi})$ denotes the marginal density of $\hat{\psi}$:

$$f\left(\hat{\psi}\right) = \int f\left(\hat{\psi}|\theta\right)p(\theta)d\theta.$$

As usual, the mode of the posterior distribution of θ can be computed by simply maximizing the value of the numerator in Eq. (87), since the denominator is not a function of θ. The marginal density of $\hat{\psi}$ is required when we want an overall measure of the fit of our model and when we want to report the shape of the posterior marginal distribution of individual elements in θ. To compute the marginal likelihood, we can use a standard random walk metropolis algorithm or a Laplace approximation. We explain the latter in Section 5.4. The results that we report are based on a standard random walk Metropolis algorithm resulting in a single Monte Carlo Markov Chain of length 600,000. The first 100,000 draws were dropped and the average acceptance rate in the chain is 27%. We confirmed that the chain is long enough so that all the statistics reported in the paper have converged. Section 6.3 compares results based on the Metropolis algorithm with the results based on the Laplace approximation.

5.3 Computation of *V*

A crucial ingredient in our empirical methodology is the matrix, $V(\theta_0, \zeta_0, T)$. The logic of our approach requires that we have at least an approximately consistent estimator of $V(\theta_0, \zeta_0, T)$. A variety of approaches are possible here. We use a bootstrap approach. Using our estimated VAR and its fitted disturbances, we generate a set of M bootstrap realizations for the impulse responses. We denote these by ψ_i, $i = 1, \ldots, M$, where ψ_i denotes the i^{th} realization of the 397×1 vector of impulse responses.[56] Consider

$$\bar{V} = \frac{1}{M}\sum_{i=1}^{M}(\psi_i - \bar{\psi})(\psi_i - \bar{\psi})', \tag{88}$$

where $\bar{\psi}$ is the mean of ψ_i, $i = 1, \ldots, M$. We set $M = 10{,}000$. The object, $\bar{V}$, is a 397 by 397 matrix, and we assume that the small sample (in the sense of T) properties of this way (or any other way) of estimating $V(\theta_0, \zeta_0, T)$ are poor. To improve small sample efficiency, we proceed in a way that is analogous to the strategy taken in the estimation of frequency-zero spectral densities (see Newey & West, 1987). In particular, rather than working with the raw variance-covariance matrix, $\bar{V}$, we instead work with $\hat{\bar{V}}$:

$$\hat{\bar{V}} = f(\bar{V}, T).$$

[56] To compute a given bootstrap realization, ψ_i, we first simulate an artificial data set, $Y_1, \ldots, Y_T$. We do this by simulating the response of our estimated VAR to an iid sequence of 14×1 shock vectors that are drawn randomly with replacement from the set of fitted shocks. We then fit a 2-lag VAR to the artificial data set using the same procedure used on the actual data. The resulting estimated VAR is then used to compute the impulse responses, which we stack into the 397×1 vector, ψ_i.

The transformation, f, has the property that it converges to the identity transform, as $T \to \infty$. In particular, $\hat{\bar{V}}$ dampens some elements in $\bar{V}$, and the dampening factor is removed as the sample grows large. The matrix, $\hat{\bar{V}}$, has on its diagonal the diagonal elements of $\bar{V}$. The entries in $\hat{\bar{V}}$ that correspond to the correlation between the l^{th} lagged response and the j^{th} lagged response in a given variable to a given shock equals the corresponding entry in $\bar{V}$, multiplied by

$$\left[1 - \frac{|l-j|}{n}\right]^{\theta_{1,T}}, l, j = 0, 1, \ldots, n.$$

Here, n denotes the number of estimated impulse response lags. Now consider the components of $\bar{V}$ that correspond to the correlations between components of different impulse response functions, either because a different variable is involved or because a different shock is involved, or both. We dampen these entries in a way that is increasing in τ, the separation in time of the two impulses. In particular, we adopt the following dampening factors for these entries:

$$\beta_T \left[1 - \frac{|\tau|}{n}\right]^{\theta_{2,T}}, \tau = 0, 1, \ldots, n.$$

We suppose that

$$\beta_T \to 1, \theta_{i,T} \to 0, \text{as } T \to \infty, i = 1, 2,$$

where the rate of convergence is whatever is required to ensure consistency of $\hat{\bar{V}}$. These conditions leave completely open what values of β_T, $\theta_{1,T}$, $\theta_{2,T}$ we use in our sample. At one extreme, we have

$$\beta_T = 0, \theta_{1,T} = \infty,$$

and $\theta_{2,T}$ unrestricted. This corresponds to the approach in CEE and ACEL, in which $\hat{\bar{V}}$ is simply a diagonal matrix composed of the diagonal components of $\bar{V}$. At the other extreme, we could set β_T, $\theta_{1,T}$, $\theta_{2,T}$ at their $T \to \infty$ values, in which $\hat{\bar{V}} = \bar{V}$. Here, we work with the approach taken in CEE and ACEL. This has the important advantage of transparency. It corresponds to selecting θ so that the model implied impulse responses lie inside a confidence tunnel around the estimated impulses. When nondiagonal terms in $\bar{V}$ are also used, then the estimator aims not just to put the model impulses inside a confidence tunnel about the point estimates, but it is also concerned about the pattern of discrepancies across different impulse responses. Precisely how the off-diagonal components of $\bar{V}$ give rise to concerns about cross-impulse response patterns of discrepancies is virtually impossible to understand intuitively. This is both because $\bar{V}$ is an enormous matrix and because it is not $\bar{V}$ that enters our criterion but its inverse.

5.4 Laplace approximation of the posterior distribution

The Metropolis algorithm for computing the posterior distribution can be time intensive, and it may be useful — at least in the intermediate stages of a research project — to use the Laplace approximation instead. In Section 6.3, we show that the two approaches generate similar results in our application, although one cannot rely on this being true in general.

To derive the Laplace approximation to $f(\theta|\hat{\psi})$, define

$$g(\theta) \equiv \log f\left(\hat{\psi}|\theta\right) + \log p(\theta).$$

Let θ^* denote the mode of the posterior distribution and define the following Hessian matrix:

$$g_{\theta\theta} = -\frac{\partial^2 g(\theta)}{\partial\theta\partial\theta'}|_{\theta=\theta^*}.$$

Note that the matrix, $g_{\theta\theta}$, is an automatic by-product of standard gradient methods for computing the mode, θ^*. The second-order Taylor series expansion of g about $\theta = \theta^*$ is

$$g(\theta) = g(\theta^*) - \frac{1}{2}(\theta - \theta^*)'g_{\theta\theta}(\theta - \theta^*),$$

where the slope term is zero if θ^* is an interior optimum, which we assume. Then,

$$f\left(\hat{\psi}|\theta\right)p(\theta) \approx f\left(\hat{\psi}|\theta^*\right)p(\theta^*)\exp\left[-\frac{1}{2}(\theta - \theta^*)g_{\theta\theta}(\theta - \theta^*)\right].$$

Note that

$$\frac{1}{(2\pi)^{\frac{m}{2}}}|g_{\theta\theta}|^{\frac{1}{2}}\exp\left[-\frac{1}{2}(\theta - \theta^*)'g_{\theta\theta}(\theta - \theta^*)\right],$$

where m denotes the number of elements in θ. The last expression is the m–variable Normal distribution for the m random variables, θ, with mean θ^* and variance-covariance matrix, $g_{\theta\theta}^{-1}$. By the standard property of a density function,

$$\int \frac{1}{(2\pi)^{\frac{m}{2}}}|g_{\theta\theta}|^{\frac{1}{2}}\exp\left[-\frac{1}{2}(\theta - \theta^*)'g_{\theta\theta}(\theta - \theta^*)\right]d\theta = 1. \tag{89}$$

Bringing together the previous results, we obtain:

$$\begin{aligned} f\left(\hat{\psi}\right) &= \int f\left(\hat{\psi}|\theta\right)p(\theta)d\theta \\ &\approx \int f\left(\hat{\psi}|\theta^*\right)p(\theta^*)\exp\left[-\frac{1}{2}(\theta - \theta^*)'g_{\theta\theta}(\theta - \theta^*)\right]d\theta \\ &= (2\pi)^{\frac{m}{2}}|g_{\theta\theta}|^{-\frac{1}{2}}f\left(\hat{\psi}|\theta^*\right)p(\theta^*), \end{aligned}$$

by Eq. (89). We now have the marginal distribution for $\hat{\psi}$. We can use this to compare the fit of different models for $\hat{\psi}$. In addition, we have an approximation to the marginal posterior distribution for an arbitrary element of θ, say θ_i :

$$\theta_i \sim N\left(\theta_i^*, \left[g_{\theta\theta}^{-1}\right]_{ii}\right),$$

where $\left[g_{\theta\theta}^{-1}\right]_{ii}$ denotes the i^{th} diagonal element of the matrix, $g_{\theta\theta}^{-1}$.

6. MEDIUM-SIZED DSGE MODEL: RESULTS

We first describe our VAR results. We then turn to the estimation of the DSGE model. Finally, we study the ability of the DSGE model to replicate the VAR-based estimates of the dynamic response of the economy to three shocks.

6.1 VAR results

We briefly describe the impulse response functions implied by the VAR. The solid line in Figures 10–12 indicate the point estimates of the impulse response functions, while the gray area displays the corresponding 95% probability bands.[57] Inflation and the interest rate are in annualized percent terms, while the other variables are measured in percent. The solid lines with squares and the dashed lines will be discussed when we review the DSGE model estimation results.

6.1.1 Monetary policy shocks

We make five observations about the estimated dynamic responses to an about 50 basis point shock to monetary policy, displayed in Figure 10. Consider first the response of inflation. Two important things to note here are the price puzzle and the delayed and gradual response of inflation.[58] In the very short run the point estimates indicate that inflation moves in a seemingly perverse direction in response to the expansionary monetary policy shock. This transitory drop in inflation in the immediate aftermath of a monetary policy shock has been widely commented upon, and has been dubbed the "price puzzle." Christiano et al. (1999) reviewed the argument that the puzzle may be the outcome of the sort of econometric specification error suggested by Sims (1992), and found evidence consistent with that view. Here, we follow ACEL and CEE in taking the position that there is no econometric specification error. Although

[57] The probability interval is defined by the point estimate of the impulse response, ±1.96 times the square root of the relevant term on the diagonal of $\bar{V}$ reported in Eq. (88).

[58] Here, we have borrowed Mankiw's (2000) language, "delayed and gradual," to characterize the nature of the response of inflation to a monetary policy shock. Although Mankiw wrote 10 years ago and he cites a wide range of evidence, Mankiw's conclusion about how inflation responds to a monetary policy shock resembles our VAR evidence very closely. Mankiw argued that the response of inflation to a monetary policy shock is gradual in the sense that it does not peak for 9 quarters.

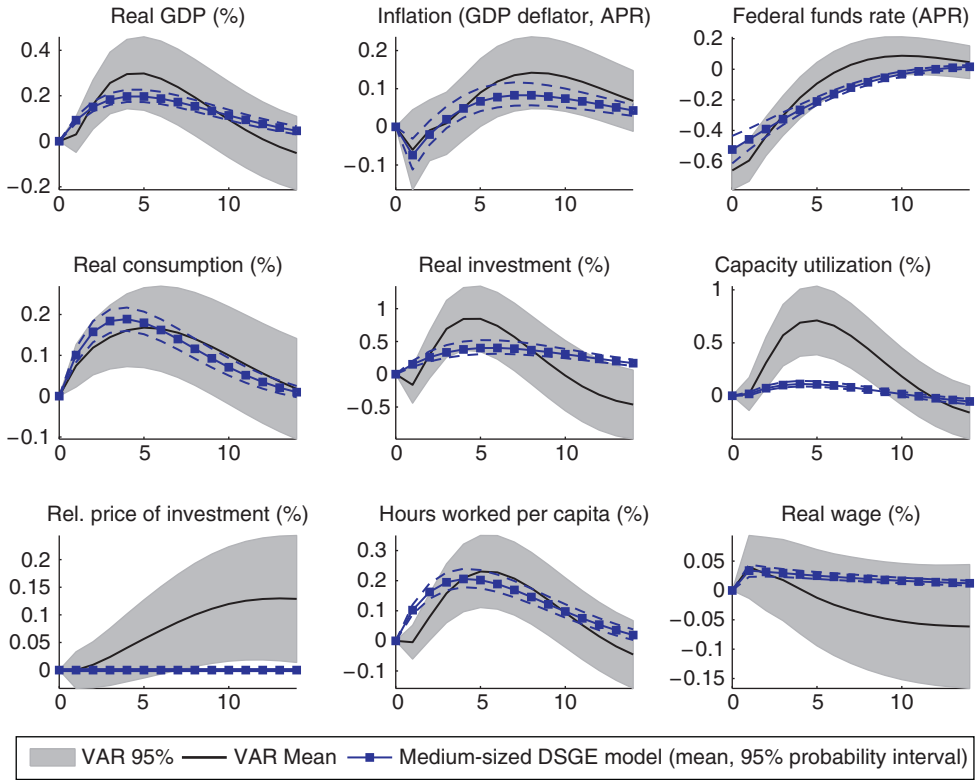

Figure 10 Dynamic responses of variables to a monetary policy shock.

the price puzzle is not statistically significant in our VAR estimation, it nevertheless deserves comment because it has potentially great economic significance. For example, the presence of a price puzzle in the data complicates the political problem associated with using high interest rates as a strategy to fight inflation. High interest rates and the consequent slowdown in economic growth are politically painful. If the public sees the high interest rate strategy producing higher inflation in the short run, support for the strategy may evaporate unless the price puzzle has been explained.[59]

Regarding the delayed and gradual response of inflation to a monetary policy shock, note how inflation reaches a peak after two years. Of course, the wide confidence

[59] There is an important historical example of this political problem. In the early 1970s, at the start of the Great Inflation in the United States, Arthur Burns was chairman of the U.S. Federal Reserve and Wright Patman was chairman of the U.S. House Committee on Banking and Currency. Patman had the opinion that, by raising costs of production, high interest rates increase inflation. Patman's belief had enormous significance because he was influential in writing the wage and price control legislation at the time. He threatened Burns that if Burns tried to raise interest rates to fight inflation, Patman would see to it that interest rates were brought under the control of the wage-price control board (see "The Lasting, Multiple Hassles of Topic A," *Time Magazine*, Monday, April 9, 1973).

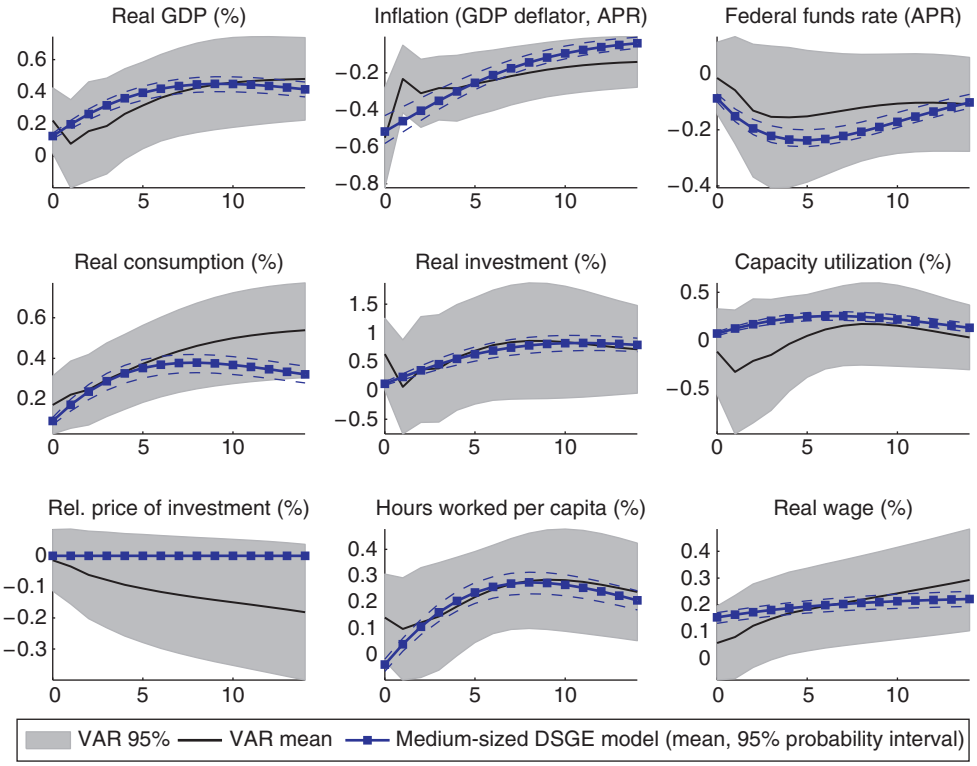

Figure 11 Dynamic responses of variables to a neutral technology shock.

intervals indicate that the exact timing of the peak is not precisely determined. However, the evidence does suggest a sluggish response of inflation. This is consistent with the views of others, arrived at by other methods, about the slow response of inflation to a monetary policy shock. As noted in the introduction to Section 4, it has been argued that the slow inflation response is a major puzzle for macroeconomics. For example, Mankiw (2000) argued that with price frictions of the type used here, the only way to explain the delayed and gradual response of inflation to a monetary policy shock is to introduce a degree of stickiness in prices that exceeds by far what can be justified based on the micro evidence. For this reason, when we study the ability of our models to match the estimated impulse response functions, we must be wary of the possibility that this is done only by making prices and wages counterfactually sticky. In addition, we must be wary of the possibility that the econometrics leans too hard on other features (such as variable capital utilization) to explain the gradual and delayed response of inflation to a monetary policy shock.

The second observation about the results in Figure 10 is that output, consumption, investment, and hours worked all display a slow, hump-shaped response to a monetary policy shock, peaking a little over one year after the shock. As emphasized in Section 4,

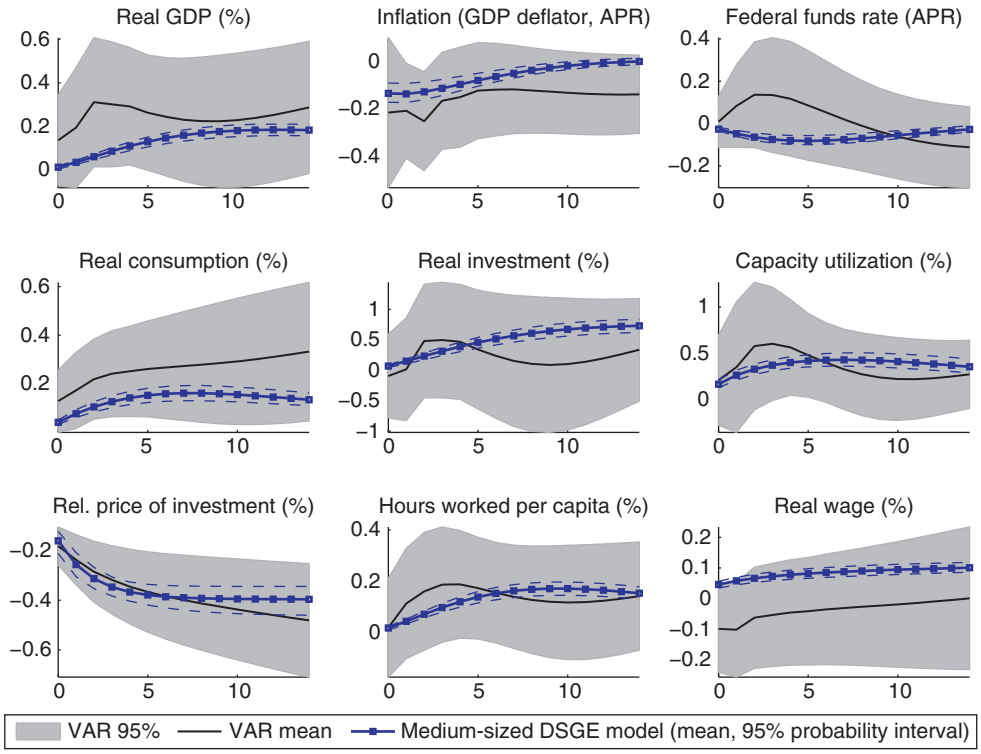

Figure 12 Dynamic responses of variables to an investment specific technology shock.

these hump-shaped observations are the reason that researchers introduce habit persistence and costs of adjustment in the flow of investment into the baseline model. Our third observation about the results in Figure 10 is that the effect of the monetary shock on the interest rate is roughly gone after two years, while the economy continues to respond well after that. This suggests that to understand the dynamic effects of a monetary policy shock, one must have a model that displays considerable sources of internal propagation.

A fourth observation concerns the response of capacity utilization. Recall from the discussion of Section 4 that the magnitude of the empirical response of this variable represents an important discipline on the analysis. In effect, those data constrain how heavily we can lean on variable capital utilization to explain the slow response of inflation to a monetary policy shock. The evidence in Figure 10 suggests that capacity utilization responds very sharply to a positive monetary policy shock. For example, utilization rises three times as much as employment, in percent terms. In interpreting this finding, we must bear in mind that the capital utilization numbers we have are for the manufacturing sector. To the extent that these data are influenced by the durable part of manufacturing, they may overstate the volatility of capacity utilization generally in the economy.

Our fifth observation about Figure 10 concerns the price of investment. In our model, this price is, by construction, unaffected by shocks other than those to the technology for converting homogeneous output into investment goods. Figure 10 indicates that the price of investment rises in response to an expansionary monetary policy shock, contrary to our model. This suggests that it would be worth exploring modifications to the technology for producing investment goods so that the trade-off between consumption and investment is nonlinear.[60] Under these conditions, the rise in the investment to consumption ratio that appears to occur in response to an expansionary monetary policy shock would be associated with an increase in the price of investment.

6.1.2 Technology shocks

Figures 11 and 12 display the responses to neutral and investment specific technology shocks, respectively. Overall, the confidence intervals are wide. The width of these confidence intervals should be no surprise in view of the nature of the question being addressed. The VAR is informed that there are two shocks in the data which have a long run effect on labor productivity, and it is being asked to determine the dynamic effects of these shocks on the data. To understand the challenge that such a question poses, imagine gazing at a data plot and thinking how the technology shocks might be detected visually. It is no wonder that in many cases, the VAR response is, "I don't know how this variable responds." This is what the wide confidence intervals tell us. For example, little can be said about the response of capacity utilization to a neutral technology shock.

Although confidence intervals are often wide, there are some responses that are significant. For example, there is a significant rise in consumption, output, and hours worked in response to a neutral shock. A particularly striking result in Figure 11 is the immediate drop in inflation in the wake of a positive shock to neutral technology. This drop has led some researchers to conjecture that the rapid response of inflation to a technology shock spells trouble for sticky price/sticky wage models. We investigate this conjecture in the next section.

6.2 Model results

6.2.1 Parameters

Parameters whose values are set a priori are listed in Table 2. We found that when we estimated the parameters κ_w and λ_w, the estimator drove them to their boundaries. This is why we simply set λ_w to a value near unity and we set $\kappa_w = 1$. The steady-state value of inflation (a parameter in the monetary policy rule and the price and wage updating equations), the steady-state government consumption to output ratio, and the steady-state growth rate of the investment specific technology were chosen to coincide with

[60] For example, instead of specifying a resource constraint in which $C_t + I_t$ appears, we could adopt one in which C_t and I_t appear in a CES function, i.e., $\left[a_1 C_t^{1/\rho} + a_2 I_t^{1/\rho}\right]^{\rho}$. The standard linear specification is a special case of this one, with $a_1 = a_2 = \rho = 1$.

Table 2 Non-Estimated Parameters in Medium-Sized DSGE Model

Parameter	Value	Description
α	0.25	Capital share
δ	0.025	Depreciation rate
β	0.999	Discount factor
π	1.0083	Gross inflation rate
η_g	0.2	Government consumption to GDP ratio
$P_{k'}$	1	Relative price of capital
κ_w	1	Wage indexation to π_{t-1}
λ_w	1.01	Wage markup
ξ_w	0.75	Wage stickiness
μ_z	1.0041	Gross neutral technology growth
μ_Ψ	1.0018	Gross investment technology growth

their corresponding sample means in our data set.[61] The growth rate of neutral technology was chosen so that, conditional on the growth rate of investment specific technology, the steady-state growth rate of output in the model coincides with the corresponding sample average in the data. We set $\xi_w = 0.75$, so that the model implies wages are reoptimized once a year on average. We did not estimate this parameter because we found that it is difficult to separately identify the value of ξ_w and the curvature parameter of household labor disutility, ϕ.

The parameters for which we report priors and posteriors are listed in Table 3. Note first that the degree of price stickiness, ξ_p, is modest. The time between price reoptimizations implied by the posterior mean of this parameter is a little less than 3 quarters. The amount of information in the likelihood, Eq. (86), about the value of ξ_p is substantial. The posterior standard deviation is roughly one-third the size of the prior standard deviation and the posterior 95% probability interval is a quarter of the width of the corresponding prior probability interval. Generally, the amount of information in the likelihood about all the parameters is large in this sense. An exception to this pattern is the coefficient on inflation in the Taylor rule, r_π. There appears to be relatively little information about this parameter in the likelihood. Note that ϕ is estimated to be quite small, implying a consumption-compensated labor supply elasticity for the household of around 8. Such a high elasticity would be regarded as empirically implausible if

[61] In our model, the relative price of investment goods represents a direct observation of the technology shock for producing investment goods.

Table 3 Prior and Posteriors of Parameters for Medium-Sized DSGE Model

		Prior		Posterior[a]
Parameter		**Distribution [bounds]**	**Mean, std. dev. [5% and 95%]**	**Mean, std. dev. [5% and 95%]**
Price-setting parameters				
Price stickiness	ξ_p	Beta [0, 0.8]	0.50, 0.15 [0.23, 0.72]	0.62, 0.04 [0.56, 0.68]
Price markup	λ_f	Gamma [1.01, ∞]	1.20, 0.15 [1.04, 1.50]	1.20, 0.08 [1.06, 1.32]
Monetary authority parameters				
Taylor rule: Interest smoothing	ρ_R	Beta [0, 1]	0.80, 0.10 [0.62, 0.94]	0.87, 0.02 [0.85, 0.90]
Taylor rule: Inflation coefficient	r_π	Gamma [1.01, 4]	1.60, 0.15 [1.38, 1.87]	1.43, 0.11 [1.25, 1.59]
Taylor rule: GDP coefficient	r_y	Gamma [0, 2]	0.20, 0.15 [0.03, 0.49]	0.07, 0.03 [0.02, 0.11]
Household parameters				
Consumption habit	b	Beta [0, 1]	0.75, 0.15 [0.47, 0.95]	0.77, 0.02 [0.74, 0.80]
Inverse labor supply elasticity	ϕ	Gamma [0, ∞]	0.30, 0.20 [0.06, 0.69]	0.12, 0.03 [0.08, 0.16]
Capacity adjustment costs curv.	σ_a	Gamma [0, ∞]	1.00, 0.75 [0.15, 2.46]	0.30, 0.08 [0.16, 0.44]
Investment adjustment costs curv.	S''	Gamma [0, ∞]	12.00, 8.00 [2.45, 27.43]	14.30, 2.92 [9.65, 18.8]
Shocks				
Autocorr. investment technology	ρ_Ψ	Uniform [0, 1]	0.50, 0.29 [0.05, 0.95]	0.60, 0.08 [0.48, 0.72]
Std. dev. neutral tech. shock (%)	σ_z	Inv. Gamma [0, ∞]	0.20, 0.10 [0.10, 0.37]	0.22, 0.02 [0.19, 0.25]
Std. dev. invest. tech. shock (%)	σ_Ψ	Inv. Gamma [0, ∞]	0.20, 0.10 [0.10, 0.37]	0.16, 0.02 [0.12, 0.20]
Std. dev. monetary shock (APR)	σ_R	Inv. Gamma [0, ∞]	0.40, 0.20 [0.21, 0.74]	0.51, 0.05 [0.44, 0.58]

[a]Based on standard random walk metropolis algorithm. 600 000 draws, 100 000 for burn-in, acceptance rate 27%.

Table 4 Medium-Sized DSGE Model Steady State at Posterior Mean for Parameters

Variable	Standard model	Description
k/γ	7.73	Capital to GDP ratio (quarterly)
c/γ	0.56	Consumption to GDP ratio
i/γ	0.24	Investment to GDP ratio
H	0.63	Steady-state labor input
R	1.014	Gross nominal interest rate (quarterly)
R^{real}	1.006	Gross real interest rate (quarterly)
r^k	0.033	Capital rental rate (quarterly)
A_L	2.22	Slope, labor disutility

it were interpreted as the elasticity of supply of hours by a representative agent. However, as discussed in Section 2.3, this is not our interpretation. Table 4 reports steady-state properties of the model, evaluated at the posterior mean of the parameters.

6.2.2 Impulse responses

We now comment on the DSGE model impulse responses displayed in Figures 10–12. The line with solid squares in the figures display the impulse responses of our model, at the posterior mean of the parameters. The dashed lines display the 95% probability interval for the impulse responses implied by the posterior distribution of the parameters. These intervals are in all cases reasonably tight, reflecting the tight posterior distribution on the parameters as well as the natural restrictions of the model.

Our estimation strategy in effect selects a model parameterization that places the model-implied impulse response functions as close as possible to the center of the gray area, while not suffering too much of a penalty from the priors. The estimation criterion is less concerned about reproducing VAR-based impulse response functions where the gray areas are the widest.

Consider Figure 10, which displays the response of standard macroeconomic variables to a monetary policy shock. Note how well the model captures the delayed and gradual response of inflation. In the model it takes two years for inflation to reach its peak response after the monetary policy shock. Importantly, the model even captures the price puzzle phenomenon, according to which inflation moves in the "wrong" direction initially. This apparently perverse initial response of inflation is interpreted by the model as reflecting the reduction in labor costs associated with the cut in the nominal rate of interest. The notable result here is that the slow response of inflation to a monetary policy shock is explained with a modest degree of wage and price-setting frictions. In addition, the gradual and delayed response of inflation is not due to an excessive or counterfactual

increase in capital utilization. Indeed, the model substantially understates the rise in capital utilization. While on its own this is a failure of the model, the weak utilization response does draw attention to the apparent ease with which the model is able to capture the inertial response of inflation to a monetary shock.

The model also captures the response of output and consumption to a monetary policy shock reasonably well. However, the model apparently does not have the flexibility to capture the relatively sharp fall and rise in the investment response, although the model responses lie inside the gray area. The relatively large estimate of the curvature in the investment adjustment cost function, S'', suggests that to allow a greater response of investment to a monetary policy shock would cause the model's prediction of investment to lie outside the gray area in the first couple of quarters. These findings for monetary policy shocks are broadly similar to those reported in CEE and ACEL.

Figure 11 displays the response of standard macroeconomic variables to a neutral technology shock. Note that the model is reasonably successful at reproducing the empirically estimated responses. The dynamic response of inflation is particularly notable, in light of the estimation results reported in ACEL. Those results suggest that the sharp and precisely estimated drop in inflation in response to a neutral technology shock is difficult to reproduce in a model like ours. In describing this problem for their model, ACEL expressed a concern that the failure reflects a deeper problem with sticky price models.[62] They suggested that perhaps the emphasis on price- and wage-setting frictions, largely motivated by the inertial response of inflation to a monetary shock, is shown to be misguided by the evidence that inflation responds rapidly to technology shocks.[63] Our results suggest a far more mundane possibility.

There are two key differences between our model and the one in ACEL that allow it to reproduce the response of inflation to a technology shock more or less exactly without hampering its ability to account for the slow response of inflation to a monetary policy shock. First, in our model there is no indexation of prices to lagged inflation (see Eq. 63). ACEL follows CEE in supposing that when firms cannot optimize their price, they index it fully to lagged aggregate inflation. The position of our model on price indexation is a key reason why we can account for the rapid fall in inflation after a neutral technical shock while ACEL cannot. We suspect that our way of treating indexation is a step in the right direction from the point of view of the microeconomic data. Micro observations suggest that individual prices do not change for extended periods of time. A second distinction between our model and the one in ACEL is that we

[62] See Paciello (2009) for another discussion of this point.

[63] The concern is reinforced by the fact that an alternative approach, one based on information imperfections and minimal price/wage-setting frictions, seems like a natural one for explaining the puzzle of the slow response of inflation to monetary policy shocks and the quick response to technology shocks (see Maćkowiak and Wiederholt, 2009; Mendes, 2009; and Paciello, 2009). Dupor, Han, and Tsai (2009) suggested more modest changes in the model structure to accommodate the inflation puzzle.

specify the neutral technology shock to be a random walk (see Eq. 60), while in ACEL the growth rate of the estimated technology shock is highly autocorrelated. In ACEL, a technology shock triggers a strong wealth effect, which stimulates a surge in demand that places upward pressure on marginal cost and thus inflation.

Figure 12 displays dynamic responses of macroeconomic variables to an investment specific shock. The DSGE model fits the dynamics implied by the VAR well, although the confidence intervals are large.

6.3 Assessing VAR robustness and accuracy of the Laplace approximation

It is well known that when the start date or number of lags for a VAR are changed, the estimated impulse response functions change. In practice, one hopes that the width of probability intervals reported in the analysis is a reasonable rule-of-thumb guide to the degree of nonrobustness. In Figures 13–15 we display all the estimated impulse response functions from our VAR when we apply a range of different start dates and lag lengths. The VAR point estimates used in our estimation exercise are displayed

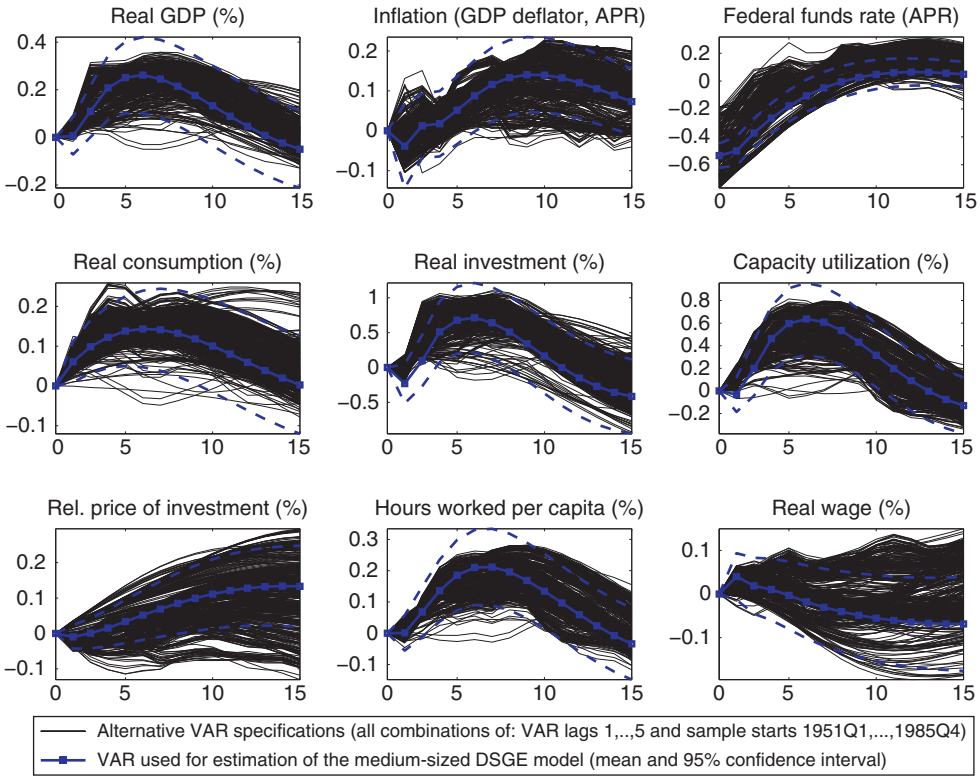

Figure 13 VAR specification sensitivity: Response to a monetary policy shock.

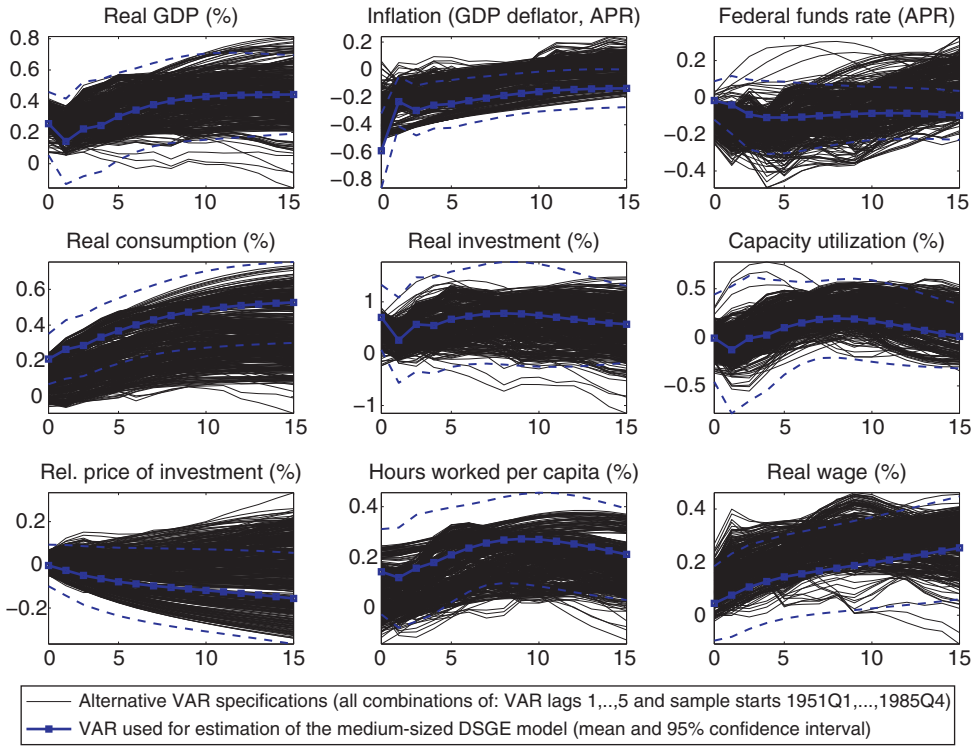

Figure 14 VAR specification sensitivity: Neutral technology shock.

in Figures 13–15 in the form of the solid line with solid squares. The 95% probability intervals associated with the impulse response functions used in our estimation exercise are indicated by the dashed lines. According to the figures, the degree of variation across different samples and lag lengths corresponds roughly to the width of probability intervals. Although results do change across the perturbed VARs, the magnitude of the changes are roughly what is predicted by the rule-of-thumb. In this sense, the degree of nonrobustness in the VAR is not great.

Finally, Figure 16 displays the priors and posteriors of the model parameters. The posteriors are computed by two methods: the random walk Metropolis method, and the Laplace approximation described in Section 5.4. It is interesting that the Laplace approximation and the results of the random walk Metropolis algorithm are very similar. These results suggest that one can save substantial amounts of time by computing the Laplace approximation during the early and intermediate phases of a research project. At the end of the project, when it is time to produce the final draft of the manuscript, one can then perform the time-intensive random walk Metropolis calculations.

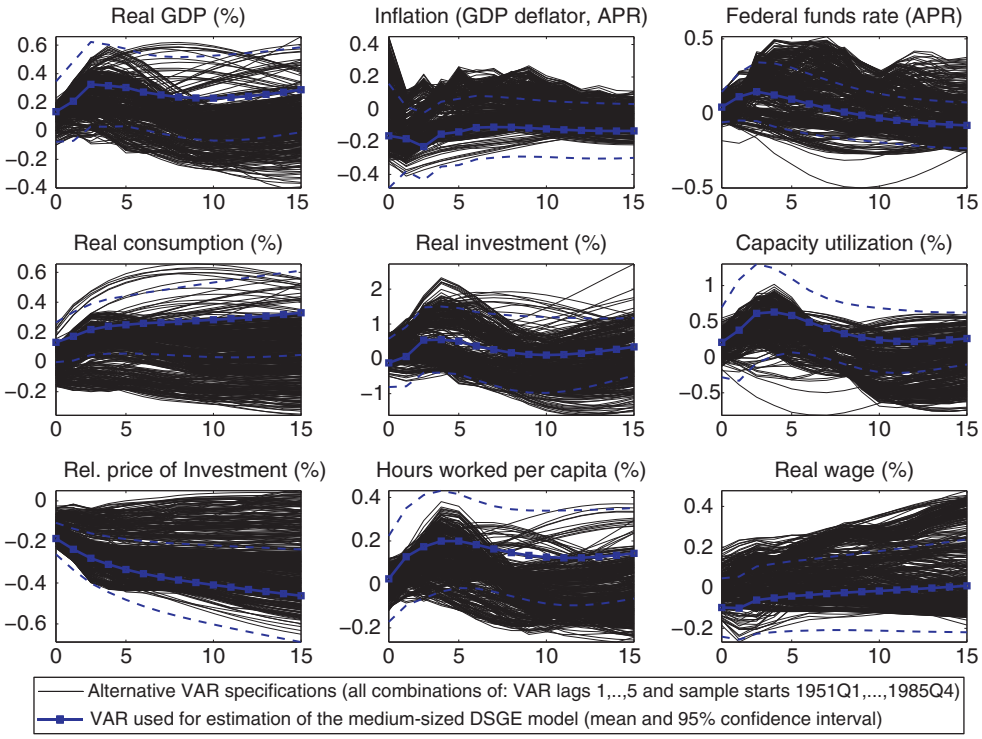

Figure 15 VAR specification sensitivity: Investment specific technology shock.

7. CONCLUSION

The literature on DSGE models for monetary policy is too large to review in all its detail in this chapter. Necessarily, we have been forced to focus on only a part. Relatively little space has been devoted to the limitations of monetary DSGE models. A key challenge is posed by the famous statistical rejections of the intertemporal Euler equation that lies at the heart of DSGE models (see, e.g., Hansen & Singleton, 1983). These rejections of the "IS equation" in the New Keynesian model pose a challenge for that model's account of the way shocks propagate through the economy. At the same time, the Bayesian impulse response matching technique that we apply suggests that the New Keynesian model is able to capture the basic features of the transmission of three important shocks.[64] An outstanding question is how to resolve these apparently conflicting pieces of information.

Also, we have been able to do little in the way of reviewing the new frontiers for monetary DSGE models. The recent financial turmoil has accelerated work to

[64] In our empirical analysis we have not reported our VAR's implications for the importance of the three shocks that we analyzed. However, ACEL documents that these shocks together account for well over 50% of the variation of macroeconomic time series like output, investment and employment.

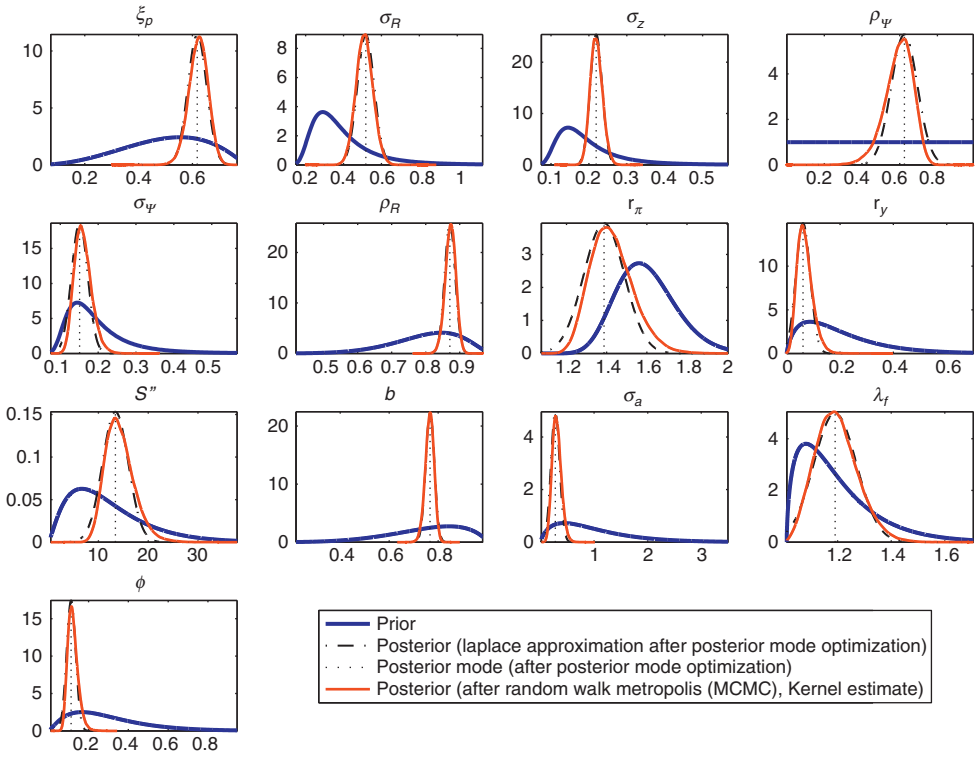

Figure 16 Priors and posteriors of estimated parameters of the medium–sized DSGE model.

introduce a richer financial sector into the New Keynesian model. With this addition, the model is able to address important policy questions that cannot be addressed by the models described here. How should monetary policy respond to an increase in interest rate spreads? How should we think about the recent "unconventional monetary policy" actions, in which the monetary authority purchases privately issued liabilities such as mortgages and commercial paper? The models described in this chapter are silent on these questions. However, an exploding literature too large to review here has begun to introduce the modifications necessary to address them.[65] The labor market is another frontier of new model development. We have presented a rough sketch of the approach in CTW, but the literature merging the best of labor market research with monetary DSGE models is too large to survey here.[66] Still, these new developments

[65] For a small sampling, see, for example, Bernanke, Gertler, and Gilchrist (1999); Christiano, Motto, and Rostagno (2003, 2009); Cúrdia and Woodford (2009); and Gertler and Kiyotaki (2010).

[66] A small open economy model with financial and labor market frictions, estimated by full information Bayesian methods, appears in Christiano, Trabandt, and Walentin (2010c). Important other papers on the integration of unemployment and other labor market frictions into monetary DSGE models include Gali (2010a); Gertler, Sala, and Trigari (2008); and Thomas (2009).

ensure that monetary DSGE models will remain an active and exciting area of research for the foreseeable future.

REFERENCES

Abel, A.B., Bernanke, B., 2005. Macroeconomics, fifth ed. Pearson Addison Wesley, Upper Saddle River, NJ.

Adalid, R., Detken, C., 2007. European Central Bank, Working Paper No. 732, Liquidity shocks and asset price boom/bust cycles. (February).

Altig, D., Christiano, L.J., Eichenbaum, M., Lindé, J., 2005. Firm-specific capital, nominal rigidities and the business cycle. NBER, Working Paper 11034.

Ball, L., 1994. Credible disinflation with staggered price setting. Am. Econ. Rev. 84 (March), 282–289.

Barth III, M.J., Ramey, V.A., 2002. The cost channel of monetary transmission. In: Bernanke, B., Rogoff, K.S. (Eds.), NBER chapters, NBER macroeconomics annual 2001. 16, MIT Press, Cambridge, MA, pp. 199–256.

Basu, S., 1995. Intermediate goods and business cycles: Implications for productivity and welfare. Am. Econ. Rev. 85 (3), 512–531.

Benhabib, J., Schmitt-Grohé, S., Uribe, M., 2002. Chaotic interest-rate rules. Am. Econ. Rev. 92 (2), 72–78.

Bernanke, B., Boivin, J., Eliasz, P.S., 2005. Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. Q. J. Econ. 120 (1), 387–422.

Bernanke, B., Gertler, M., 2000. Monetary policy and asset price volatility. NBER, Working Paper No. 7559.

Bernanke, B., Gertler, M., Gilchrist, S., 1999. The financial accelerator in a quantitative business cycle framework. In: Taylor, J.B., Woodford, M. (Eds.), Handbook of macroeconomics. Elsevier Science, North-Holland, Amsterdam, pp. 1341–1393.

Boldrin, M., Christiano, L.J., Fisher, J.D.M., 2001. Habit persistence, asset returns, and the business cycle. Am. Econ. Rev. 91 (1), 149–166.

Bruckner, M., Schabert, A., 2003. Supply-side effects of monetary policy and equilibrium multiplicity. Econ. Lett. 79 (2), 205–211.

Bullard, J., Mitra, K., 2002. Learning about Monetary policy rules. J. Monetary Econ. 49, 1105–1129.

Calvo, G.A., 1983. Staggered prices in a utility-maximizing framework. J. Monetary Econ. 12 (3), 383–398.

Carroll, C.D., Overland, J., Weil, D.N., 1997. Comparison utility in a growth model. J. Econ. Growth 2 (4), 339–367.

Carroll, C.D., Overland, J., Weil, D.N., 2000. Saving and growth with habit formation. Am. Econ. Rev. 90 (3), 341–355.

Chernozhukov, V., Hong, H., 2003. An MCMC approach to classical estimation. J. Econom. 115 (2), 293–346.

Chowdhury, I., Hoffmann, M., Schabert, A., 2006. Inflation dynamics and the cost channel of monetary transmission. Eur. Econ. Rev. 50 (4), 995–1016.

Christiano, L.J., 1988. Why does inventory investment fluctuate so much? J. Monetary Econ. 21, 247–280.

Christiano, L.J., 1991. Modeling the liquidity effect of a money shocks. Federal Reserve Bank of Minneapolis Quarterly Review 15 (1), 1–34.

Christiano, L.J., 2004. Firm-specific capital and aggregate inflation dynamics in Woodford's model. Manuscript.

Christiano, L.J., 2007. Discussion of Del Negro, Schorfheide, Smets and Wouters. J. Bus. Econ. Stat. 25 (2), 143–151.

Christiano, L.J., Eichenbaum, M., 1992a. Current real business cycle theories and aggregate labor market fluctuations. Am. Econ. Rev. 82 (3), 430–450.

Christiano, L.J., Eichenbaum, M., 1992b. Liquidity effects and the monetary transmission mechanism. American Economic Review, Papers and Proceedings 82, 346–353.

Christiano, L.J., Eichenbaum, M., Evans, C.L., 1996, The effects of monetary policy shocks: evidence from the flow of funds, The Review of Economics and Statistics, Vol. 78, No. 1 (February), pp. 16–34.
Christiano, L.J., Eichenbaum, M., Evans, C.L., 1999. Monetary policy shocks: What have we learned and to what end? In: Taylor, J.B., Woodford, M. (Eds.), Handbook of macroeconomics. Elsevier Science, North-Holland, Amsterdam, pp. 65–148.
Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. J. Polit. Econ. 113 (1), 1–45.
Christiano, L.J., Gust, C., 2000. The expectations trap hypothesis. Federal Reserve Bank of Chicago Economic Perspectives 24, 21–39.
Christiano, L.J., Ilut, C., Motto, R., Rostagno, M., 2008. Monetary policy and stock market boom bust cycles. European Central Bank working paper number 955.
Christiano, L.J., Ilut, C., Motto, R., Rostagno, M., 2010, Monetary Policy and Stock Market Booms, paper presented to conference sponsored by Federal Reserve Bank of Kansas City, 'Macroeconomic Challenges: the Decade Ahead,' at Jackson Hole, Wyoming, August 26–28.
Christiano, L.J., Motto, R., Rostagno, M., 2003. The great depression and the Friedman-Schwartz hypothesis. J. Money Credit Bank. 35 (6) December, Part 2, 1119–1197.
Christiano, L.J., Motto, R., Rostagno, M., 2009. Financial factors in business cycles. Manuscript.
Christiano, L.J., Rostagno, M., 2001. Money growth monitoring and the Taylor rule. NBER, Working Paper, No. 8539.
Christiano, L.J., Trabandt, M., Walentin, K., 2010a. Involuntary unemployment and the business cycle. NBER, Working Paper No. 15801.
Christiano, L.J., Trabandt, M., Walentin, K., 2010b. A monetary business cycle model with labor market frictions. Northwestern University Manuscript.
Christiano, L.J., Trabandt, M., Walentin, K., 2010c. Introducing financial frictions and unemployment into a small open economy model. Sveriges Riksbank, Working Paper No. 214.
Clarida, R., Gali, J., Gertler, M., 1999. The science of monetary policy: A New Keynesian perspective. J. Econ. Lit. 37, 1661–1707.
Cochrane, J.H., 2009. Can learnability save New-Keynesian models? J. Monetary Econ. 56, 1109–1113.
Constantinides, G.M., 1990. Habit formation: A resolution of the equity premium puzzle. J. Polit. Econ. 98 (3), 519–543.
Cúrdia, V., Woodford, M., 2009. Credit spreads and monetary policy. NBER, Working Paper No. 15289.
de Walque, G., Smets, F., Wouters, R., 2006. Firm-specific production factors in a DSGE model with Taylor price setting. International Journal of Central Banking 2 (3), 107–154.
Dupor, B., Han, J., Tsai, Y.C., 2009. What do technology shocks tell us about the New Keynesian paradigm? J. Monetary Econ. 56 (4), 560–569.
Erceg, C.J., Henderson, D.W., Levin, A.T., 2000. Optimal monetary policy with staggered wage and price contracts. J. Monetary Econ. 46, 281–313.
Fernald, J., 2009. A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco Manuscript.
Fisher, J., 2006. The dynamic effects of neutral and investment-specific technology shocks. J. Polit. Econ. 114 (3), 413–451.
Fuerst, T.S., 1992. Liquidity, loanable funds, and real activity. J. Monetary Econ. 29 (1), 3–24.
Fuhrer, J.C., 2000. Habit formation in consumption and its implications for monetary policy models. Am. Econ. Rev. 90 (3), 367–390.
Gali, J., 1999. Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? Am. Econ. Rev. 89 (1), 249–271.
Gali, J., 2010a. The return of the wage Phillips curve. CREI Manuscript.
Gali, J., 2010b. Unemployment fluctuations and stabilization policies: A New Keynesian perspective. Department of Economics, University of Copenhagen Zeuthen lectures delivered March 19–20 to Zeuthen Workshop on Macroeconomics.
Gertler, M., Kiyotaki, N., 2010. Financial intermediation and credit policy in business cycle analysis. Manuscript.

Gertler, M., Sala, L., Trigari, A., 2008. An estimated monetary DSGE model with unemployment and staggered nominal wage bargaining. J. Money Credit Bank. 40 (8), 1713–1764.
Guerron-Quintana, P., 2008. Refinements on macroeconomic modeling: The role of non-separability and heterogeneous labor supply. J. Econ. Dyn. Control 32, 3613–3630.
Hamilton, J., 1994. Time series analysis. Princeton University Press, Princeton, NJ.
Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. Econometrica 50, 1029–1054.
Hansen, L.P., Singleton, K.J., 1983. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. J. Polit. Econ. 91 (2), 249–265.
Hansen, G.D., 1985. Indivisible labor and the business cycle. J. Monetary Econ. 16 (3), 309–327.
Hodrick, R.J., Prescott, E.C., 1997. Postwar U.S. business cycles: An empirical investigation. J. Money Credit Bank. 29 (1), 1–16 February.
Justiniano, A., Primiceri, G., 2008. Potential and natural output. Northwestern University Manuscript.
Kiley, M.T., 2010. Output gaps. Federal Reserve Board of Governors Finance and Economics Discussion Series, Working Paper 2010–27.
Kim, J.Y., 2002. Limited information likelihood and Bayesian analysis. J. Econom. 107, 175–193.
King, R.G., Rebelo, S., 1993. Low frequency filtering and real business cycles. J. Econ. Dyn. Control 17 (1–2), 207–231.
Krusell, P., Mukoyama, T., Rogerson, R., Sahin, A., 2008. Aggregate implications of indivisible labor, incomplete markets and labor market frictions. NBER, Working Paper No. 13871.
Kwan, Y.K., 1999. Asymptotic Bayesian analysis based on a limited information estimator. J. Econom. 88, 99–121.
Levin, A.T., Onatski, A., Williams, J.C., Williams, N., 2005. Monetary policy under uncertainty in micro-founded macroeconometric models. NBER Macroeconomics Annual 20, 229–287.
Lucas Jr., R.E., Prescott, E.C., 1971. Investment under uncertainty. Econometrica 39 (5), 659–681.
Lucca, D.O., 2006. Essays in investment and macroeconomics. Northwestern University, Department of Economics, Ph.D. dissertation.
Maćkowiak, B., Wiederholt, M., 2009. Optimal sticky prices under rational inattention. Am. Econ. Rev. 99 (3), 769–803.
Mankiw, N.G., 2000. The inexorable and mysterious tradeoff between inflation and unemployment. NBER, Working Paper 7884.
Matsuyama, K., 1984. A learning effect model of investment: An alternative interpretation of Tobin's q. Northwestern University, Manuscript.
McCallum, B.T., 2009. Inflation determination with Taylor rules: Is New-Keynesian analysis critically flawed? J. Monetary Econ. 56, 1101–1108.
Mendes, R., 2009. Information, expectations, and the business cycle. Research Department, Bank of Canada Manuscript.
Mulligan, C.B., 2001. Aggregate implications of indivisible labor. Advances in Macroeconomics 1 (1) Article 4, Berkeley Electronic Press, pp. 1–33.
Newey, W.K., West, K.D., 1987. A simple, positive semi-definite heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55 (3), 703–708.
Paciello, L., 2009. Does inflation adjust faster to aggregate technology shocks than to monetary policy shocks?. Einaudi Institute for Economics and Finance, Working Paper.
Prescott, E.C., 1986. Theory ahead of business-cycle measurement. Carnegie-Rochester Conference Series on Public Policy 25 (1), 11–44.
Prescott, E.C., 1998. Needed: A theory of total factor productivity. Int. Econ. Rev. 39 (3), 525–551.
Ravenna, F., Walsh, C.E., 2006. Optimal monetary policy with the cost channel. J. Monetary Econ. 53, 199–216.
Rogerson, R., 1988. Indivisible labor, lotteries and equilibrium. J. Monetary Econ. 21 (1), 3–16.
Rogerson, R., Wallenius, J., 2009. Micro and macro elasticities in a life cycle model with taxes. J. Econ. Theory 144, 2277–2292.

Rotemberg, J., Woodford, M., 1997. An optimization-based econometric framework for the evaluation of monetary policy. In: Bernanke, B., Rotemberg, J. (Eds.), NBER macroeconomics annual. MIT Press, Cambridge, MA.

Sala, L., Söderström, U., Trigari, A., 2008. Monetary policy under uncertainty in an estimated model with labor market frictions. J. Monetary Econ. 55 (5), 983–1006.

Sargent, T., 1979. Macroeconomic theory. Academic Press, New York.

Sims, C.A., 1992. Interpreting the macroeconomic time series facts: The effects of monetary policy. Eur. Econ. Rev. 36 (5), 975–1000.

Sims, C.A., Zha, T., 2006. Were there regime switches in U.S. monetary policy? Am. Econ. Rev. 96 (1), 54–81.

Smets, F., Wouters, R., 2007. Shocks and frictions in US business cycles. Am. Econ. Rev. 97 (3), 586–606.

Sveen, T., Weinke, L., 2005. New perspectives on capital, sticky prices, and the Taylor principle. J. Econ. Theory 123 (1), 21–39.

Thomas, C., 2009. Search frictions, real rigidities and inflation dynamics. Bank of Spain Manuscript.

Topel, R., Rosen, S., 1988. Housing investment in the United States. J. Polit. Econ. 96 (4), 718–740.

Walsh, C., 2005. Discussion of Levin, Onatski, Williams and Williams. 20, MIT Press, Cambridge, MA NBER macroeconomics annual.

Woodford, M., 2003. Interest and prices: Foundations of a theory of monetary policy. Princeton University Press, Princeton, NJ.

Woodford, M., 2004. Inflation and output dynamics with firm-specific investment. Manuscript.

Yun, T., 1996. Nominal price rigidity, money supply endogeneity, and business cycles. J. Monetary Econ. 37 (2), 345–370.