

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

**The Frequency of Price Adjustment and New
Keynesian Business Cycle Dynamics**

Richard Dennis
Federal Reserve Bank of San Francisco

November 2007

Working Paper 2006-22

<http://www.frbsf.org/publications/economics/papers/2006/wp06-22bk.pdf>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System. This paper was produced under the auspices for the Center for Pacific Basin Studies within the Economic Research Department of the Federal Reserve Bank of San Francisco.

The Frequency of Price Adjustment and New Keynesian Business Cycle Dynamics*

Richard Dennis[†]
Federal Reserve Bank of San Francisco

November 28, 2007

Abstract

The Calvo pricing model that lies at the heart of many New Keynesian business cycle models has been roundly criticized for being inconsistent both with time series data on inflation and with micro-data on the frequency of price changes. In this paper I develop a new pricing model whose structure can be interpreted in terms of menu costs and information gathering/processing costs, that usefully addresses both criticisms. The resulting Phillips curve encompasses the partial-indexation model, the full-indexation model, and the Calvo model, and can speak to micro-data in ways that these models cannot. Taking the menu/information costs Phillips curve to the data, I find that the share of firms that change prices each quarter is about 60 percent and, reflecting the importance of information gathering/processing costs, that most firms that change prices are rule-of-thumb price setters. Exploiting an isomorphism result, I show that these values are consistent with estimates implied by the partial-indexation model.

Keywords: *Price adjustment, inflation indexation, Bayesian estimation.*

JEL Classification: C11, C52, E31, E52.

*I would like to thank John Fernald and seminar participants at Bocconi University, the Federal Reserve Bank of San Francisco, Tokyo University, the Reserve Bank of New Zealand, Boston College, and the Royal Economic Society conference, 2007, for comments. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

[†]Address for Correspondence: Economic Research, Mail Stop 1130, Federal Reserve Bank of San Francisco, 101 Market St, CA 94105, USA. Email: richard.dennis@sf.frb.org.

1 Introduction

New Keynesian business cycle models have become the dominant framework for studying the design and conduct of monetary policy. The models formalize the rigidities and market imperfections that govern their behavior and are micro-founded, which permits welfare analysis and makes policy experiments conducted within them immune to the Lucas (1976) critique. Prominent examples in the New Keynesian tradition include Rotemberg and Woodford (1997), Clarida, Gali, and Gertler (1999), McCallum and Nelson (1999), Walsh (2003), and Woodford (2003). One of the most important components in these models is the New Keynesian Phillips curve, the equation linking inflation to marginal costs that provides a stabilization role for monetary policy. The micro-structure that is most widely used to derive the New Keynesian Phillips curve is the Calvo model¹ (Calvo, 1983), and the defining feature of this model is that only a fixed (Calvo-) share of firms have the opportunity to optimize their prices each period. This Calvo-share parameter governs the frequency with which firms change prices and determines the average duration between price changes.

Despite its popularity, the New Keynesian Phillips curve has attracted considerable criticism. Some criticisms are empirical; Estrella and Fuhrer (2002) argue that the New Keynesian Phillips curve provides a poor description of inflation dynamics because it asserts a correlation structure among inflation, the change in inflation, and marginal costs that prevents it from replicating the hump-shaped responses that are widely recognized to characterize inflation's behavior following shocks.² Similarly, Rudd and Whelan (2006) argue that the New Keynesian Phillips curve is incapable of describing inflation dynamics and suggest that there is little evidence of the type of forward-looking behavior required by the model. Other criticisms focus on whether estimates of the New Keynesian Phillips curve are economically plausible. In this vein, a prominent criticism is that Calvo-shares estimated from the New Keynesian Phillips curve imply a level of price rigidity that is inconsistent with micro-data on the frequency of price adjustment. For example, Eichenbaum and Fisher (2007) estimate the Calvo-share to be around 0.85 for the United States, which implies that only 15 percent of firms change their prices each quarter and that firms change prices once every 20 months on average. But after examining Bureau of Labor Statistics data on price changes — the very price data that go into

¹Roberts (1995) shows that Rotemberg's (1982) quadratic price adjustment costs model and Taylor's (1980) overlapping nominal wage contracts model give rise to closely related specifications, so the issues discussed in this paper apply equally to these models.

²In fact, the Estrella and Fuhrer (2002) criticisms apply to an entire class of rational expectations models, not just to the New Keynesian Phillips curve.

the consumer price index and the personal consumption expenditures price index — Bils and Klenow (2004) and Nakamura and Steinsson (2006) report that, excluding temporary sales, the average duration between price changes for the expenditure-weighted median good is 5.5 months and 8.6 months, respectively. The disparity between estimates of the Calvo-share and micro-evidence on the frequency of price adjustment is worrisome, particularly since models built around the New Keynesian Phillips curve are routinely used to address issues as important as how to design a welfare-maximizing monetary policy (Erceg, Henderson, and Levin, 2000).

In this paper, I develop a new model of price setting and show that it usefully addresses these criticisms of the New Keynesian Phillips curve. An essential feature of this model is that, while a share of firms have the opportunity to change prices each period, they do not necessarily make an optimal price change. Instead, among those firms that change prices a fraction makes an optimal price change, while the remainder employ an indexation pricing strategy. In this respect (but not in others) the pricing model is similar to Galí and Gertler (1999). Why is this price-setting environment attractive? Where traditional models of price adjustment have emphasized physical costs to changing prices, such as menu costs, as the source of price rigidity (Mankiw, 1985), recent literature has emphasized the costs that firms face when gathering (Mankiw and Reis, 2002) and processing (Sims, 2003) the information they require in order to set prices optimally. In fact, evidence suggests that costs to gathering and processing information and company managerial and organizational issues (Zbaraki, Ritson, Levy, Dutta, and Bergin, 2004; Zbaraki, Levy, and Bergin, 2007) may be much more important for price setting than traditional menu cost factors. An attractive aspect of the price-setting environment that I analyze, then, is that it provides a vehicle through which both costs can play a role. Menu costs — which are incurred whether or not a price change is optimal — are associated with the share of firms that can change prices. When these menu costs are large, a smaller share of firms will change their prices. Similarly, costs to gathering and processing information are associated with the share of price changers that use price indexation. When the costs to gathering and processing information are high, a larger share of price-changing firms will resort to an indexation-based pricing strategy. In addition, unlike the pricing model developed by Christiano, Eichenbaum, and Evans (2005) that asserts, contrary to the data, that all firms change price every period, the model that I develop has three distinct pricing states and allows some prices to remain unchanged between periods. As a consequence, the model's structure allows it to speak to micro-data in ways that the Christiano, Eichenbaum,

and Evans (2005) model cannot. Moreover, since some firms index their prices to past inflation, the model contains a mechanism to generate intrinsic inflation persistence.

After describing the menu/information costs model, I derive its associated Phillips curve, highlighting its connections to the New Keynesian Phillips curve and to the full- and partial-indexation Phillips curves. Specifically, I show that these alternatives are all special cases of the Phillips curve I derive and argue that the menu/information costs model's micro-structure makes it superior to these alternatives as a consequence. Subsequently, I develop a small-scale New Keynesian business cycle model and estimate specifications based on the menu/information costs Phillips curve, the Calvo Phillips curve, the full-indexation Phillips curve, and the Galí-Gertler Phillips curve. The specifications are estimated using both full information maximum likelihood (FIML) and Bayesian methods, with Bayesian predictive densities and posterior model probabilities employed for model comparison purposes.

The results are striking. First, whereas estimates of the New Keynesian Phillips curve imply an average duration between price changes that is clearly inconsistent with Bureau of Labor Statistics price data, the menu/information costs model does much better. In fact, my results place the share of firms that change prices each quarter at about 60 percent, suggesting relatively frequent price adjustment. Second, although firms change prices frequently, I find that the majority of these firms use price indexation, supporting the view that information gathering/processing costs are more important for price setting than traditional menu costs. Third, by constructing predictive densities and using Bayesian model averaging, I quantify the economy's response to technology shocks, monetary policy shocks, and consumption preference shocks, revealing the counterfactual behavior of the Calvo model and establishing that the menu/information costs model generates the hump-shaped impulse responses widely accepted to characterize the behavior of output and inflation in the U.S.

I begin by describing the New Keynesian Phillips curve and illustrating the empirical disparity between the Calvo-share and the frequency of price adjustment implied by micro-data. Section 3 outlines the economic environment that underlies the menu/information costs model and derives the associated Phillips curve. Section 4 compares the menu/information costs model to the Calvo model, the full-indexation model, and the partial-indexation model and proves its isomorphism with the partial-indexation model. Section 5 develops a small-scale New Keynesian business cycle model suitable for estimation, describes the data, and discusses the estimation strategy. Section 6 presents and interprets the estimates and compares them to those obtained from alternative pricing models. Section 7 constructs predictive densities

and uses Bayesian model averaging to summarize how consumption, inflation, and interest rates respond to shocks. Section 8 concludes.

2 The New Keynesian Phillips curve and price rigidity

As noted in the introduction, the centerpiece to much business cycle and policy analysis is the New Keynesian Phillips curve

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \frac{(1-\xi)(1-\beta\xi)}{\xi} \widehat{mc}_t, \quad (1)$$

where $\hat{\pi}_t$ and \widehat{mc}_t represent the percentage point deviation of inflation, π_t , and the percent deviation of real marginal costs, mc_t , around their zero-inflation nonstochastic steady state values, respectively. The economic environment that gives rise to this Phillips curve is one in which firms are monopolistically competitive, renting capital and labor and setting their prices to maximize profits subject to a constant elasticity of substitution demand curve, a Cobb-Douglas production technology, and a price rigidity, á la Calvo (1983). In equation (1), $\beta \in (0, 1)$ is the subjective discount factor and $\xi \in (0, 1)$ is the Calvo-share, the share of firms that cannot optimize their prices each period.

With regard to suitable values for ξ , a touchstone in the literature is Blinder (1994), who surveyed firms on the frequency of their price changes. Based on Blinder's (1994) survey, Rotemberg and Woodford (1997) set $\xi = 0.66$, which implies an average duration between price changes of nine months. But many calibration studies have assumed that prices change somewhat less frequently than this. For example, Erceg, Henderson, and Levin (2000) and Liu and Phaneuf (2005) each set $\xi = 0.75$, implying an average duration between price changes of 12 months.

Among studies that estimate ξ , a popular approach is to apply a generalized method of moments estimator to the moment condition³

$$E_t \left[\left(\hat{\pi}_t - \beta \hat{\pi}_{t+1} - \frac{(1-\xi)(1-\beta\xi)}{\xi} \widehat{mc}_t \right) \mathbf{z}_t \right] = 0, \quad (2)$$

where \mathbf{z}_t is a vector containing econometric instruments. This is the approach taken by Galí and Gertler (1999), Galí, Gertler, and López-Salido (2001), Eichenbaum and Fisher (2004), Jung and Yun (2005), and Ravenna and Walsh (2006). An alternative method is to iterate forward over equation (1) and combine the result with an evolution process for real marginal

³An alternative moment condition that is often used is equation (2) multiplied through by ξ . Some of the estimates shown in Table 1 come from this alternative moment condition.

costs to produce an estimable expression relating inflation to real marginal costs (Sbordone, 2002). A range of estimates of ξ for the U.S. are displayed in Table 1.⁴

Study	Sample	ξ
Galí & Gertler (1999)	1960:1 – 1997:4	0.829 – 0.884
Galí, Gertler & López-Salido (2001)	1970:1 – 1998:4	0.845 – 0.867
Sbordone (2002)	1960:2 – 1997:1	0.792
Eichenbaum & Fisher (2004)	1959:1 – 2001:4	0.87 – 0.91
Jung & Yun (2005)	1967:1 – 2004:4	0.910
Ravenna & Walsh (2006)	1960:1 – 2001:1	0.758 – 0.911

The estimates of ξ shown in Table 1 vary from a low of 0.758 to a high of 0.911. While $\xi = 0.758$ is broadly on par with the value used in calibration studies, a value such as $\xi = 0.911$ is much larger than either the values used in calibration exercises or the value implied by Blinder’s (1994) study. The average value for ξ in Table 1 is in the order of 0.85, suggesting that firms only change prices once every 20 months. The estimates in Table 1 highlight what has become an important criticism of the New Keynesian Phillips curve, which is that estimates of ξ are too large, implying average durations between price changes that are inconsistent with micro-evidence on the frequency of price adjustment (Bils and Klenow, 2005; Nakamura and Steinsson, 2006).⁵

2.1 Strategic complementarity and firm-specific capital

One way to resolve the apparent inconsistency between macro- and micro-estimates of the frequency of price adjustment is to change the New Keynesian Phillips curve’s micro-foundations to allow for factors such as strategic complementarity (Woodford, 2003) and/or firm-specific capital (Sbordone, 2002). These changes add one or more structural parameters to the coefficient on real marginal costs, thereby permitting greater flexibility with respect to the choice of ξ . Unfortunately, because these modifications leave the Phillips curve’s structure unchanged, they cannot, in isolation, overcome the criticism that the New Keynesian Phillips curve pro-

⁴All of the estimates reported in Table 1 have been made consistent with a Cobb-Douglas production technology and rental markets for capital and labor, facilitating comparison across studies by making the estimates invariant to particular assumptions about the steady state markup and labor’s share of income. However, the values shown may differ from those reported in the original papers as a consequence. With respect to Sbordone’s estimates, the best-fitting specification in Sbordone (2002, Table 2) has a coefficient on real marginal costs equaling $\frac{1}{18.3}$. Using Sbordone’s assumption about the discount factor and assuming a rental market for capital, the implied value for ξ is 0.792.

⁵Of course, there are other notable studies that look at micro-data on the frequency of price adjustment, including Cecchetti (1986), Carlton (1986), and Kashyap (1995).

vides a poor description of inflation dynamics (Estrella and Fuhrer, 2002; Rudd and Whelan, 2006).

3 A menu/information costs model

Firms are assumed to be monopolistically competitive and to produce according to a constant-returns-to-scale production technology subject to a downward-sloping demand schedule. In the spirit of Calvo (1983), not all firms can change their prices each period, and, in the spirit of Christiano, Eichenbaum, and Evans (2005), not all price changes that do occur are chosen optimally. However, unlike Calvo (1983), in which firms either set their prices optimally or keep their prices unchanged, and unlike Christiano, Eichenbaum, and Evans (2005), in which firms either set their prices optimally or index their prices to past inflation, in the model developed here firms are randomly allocated among three pricing states. Depending on draws from two independent Bernoulli distributions, a firm either sets its price optimally, sets its price using an indexation rule, or keeps its price unchanged. Informally, the parameters that govern the share of firms allocated to each pricing state are interpreted in terms of menu costs and the costs associated with gathering and processing the information needed to set prices optimally. For this reason, I refer to this model as the menu/information costs model.

The menu/information costs model is also related to one developed by Galí and Gertler (1999); the two models share the three distinct pricing strategies outlined above. However, there are several important differences between the two models. One important difference is that firms in the menu/information costs model internalize the three pricing strategies when optimizing their price. In contrast, optimizing firms in the Galí-Gertler model behave like those in the Calvo (1983) model. Thus, where all firms are identical and are allocated randomly among pricing states in the menu/information costs model, in the Galí-Gertler model there are two distinct types of firm: Calvo price-setters and rule-of-thumb price-setters. Another difference is that where the Galí-Gertler model contains rule-of-thumb price-setters the menu/information costs model is built around price indexation.

3.1 Basic structure

The economy is populated by a unit-measure continuum of monopolistically competitive firms. The i 'th firm, $i \in [0, 1]$, produces its differentiated product according to the Cobb-Douglas production technology

$$y_t(i) = [e^{u_t} l_t(i)]^\kappa k_t(i)^{1-\kappa}, \quad (3)$$

$\kappa \in (0, 1)$, where e^{u_t} is an aggregate labor-augmenting technology shock and $y_t(i)$, $l_t(i)$, and $k_t(i)$ denote the i 'th firm's output, labor, and capital, respectively. Firms rent capital and hire labor in perfectly competitive markets and, because they face identical factor prices, employ capital and labor in the same ratio and share the same real marginal cost, i.e., $mc_t(i) = mc_t$, $\forall i \in [0, 1]$.

A final good, Y_t , is produced from the outputs of the monopolistically competitive firms according to the Dixit and Stiglitz (1977) constant-returns-to-scale production technology

$$Y_t = \left[\int_0^1 y_t(i)^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}}, \quad (4)$$

where $\epsilon \in (1, \infty)$ is the elasticity of substitution between intermediate goods. Final goods are used for consumption and investment and are sold to households in a perfectly competitive market. Efficient production of the final good implies that the demand schedule for the i 'th firm's output takes the form

$$y_t(i) = Y_t \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon}, \quad (5)$$

where $P_t(i)$ is the price charged by the i 'th firm and P_t is the aggregate price index, the price of the final good.

Each period a fixed proportion of firms, $1 - \theta$, $\theta \in [0, 1)$, are able to change prices. However, not all firms that change prices do so optimally. Within the share of firms that change prices, a fixed proportion, $1 - \omega$, $\omega \in [0, 1)$, change their prices optimally, while the remaining proportion, ω , set their prices using the indexation rule

$$P_t(i) = (1 + \pi_{t-1}) P_{t-1}(i), \quad (6)$$

where π_t denotes the inflation rate of the final good. Unlike the Calvo model, in which firms either set their prices optimally or keep their prices unchanged, here firms are distributed among three pricing states. Specifically, each period a measure equaling θ of firms do not change their prices, a measure equaling $\omega(1 - \theta)$ of firms change their prices using the indexation rule, and a measure equaling $(1 - \omega)(1 - \theta)$ of firms set their prices to maximize expected discounted profits, with firms falling randomly into one of these three pricing states independently of their history of price changes.

To interpret this pricing structure, note that θ and ω can each be associated with a distinct cost impinging on the firm's pricing decision. The first set of costs, menu costs, are borne by firms when they change prices, regardless of whether the price change is optimal or not; these

costs are associated with θ . The second set of costs are those connected to the information gathering (Mankiw and Reis, 2002) and information processing (Sims, 2003) needed to determine the optimal price; these costs are associated with ω . Importantly, θ , which represents a cost to changing prices, determines the share of firms that change prices, not the share of firms that set their prices optimally. Because estimates of the frequency of price adjustment obtained from micro-data, such as Bils and Klenow (2004) and Nakamura and Steinsson (2006), are based on observed price changes, their findings are best interpreted as estimates of θ , the proportion of firms that change prices, rather than as estimates of $(1 - \omega)(1 - \theta)$, the proportion of firms that change prices optimally.⁶

The menu/information costs Phillips curve derived below is obtained by approximating the model around a zero-inflation steady state. In Appendix B, I treat the more general case in which the approximation is taken around a non-zero-inflation steady state. There I show that plausible values for steady state inflation do not have a large effect on the Phillips curve's coefficients.

3.2 The Dixit-Stiglitz aggregate price

With the indexation rule given by equation (6), I show in Appendix A that the aggregate price, the price of the final good, equals

$$\begin{aligned} P_t &\equiv \left[\int_0^1 P_t(i)^{1-\epsilon} di \right]^{\frac{1}{1-\epsilon}}, \\ &= \left[(1 - \theta)(1 - \omega) P_t^{*1-\epsilon} + \omega(1 - \theta)(1 + \pi_{t-1})^{1-\epsilon} P_{t-1}^{1-\epsilon} + \theta P_{t-1}^{1-\epsilon} \right]^{\frac{1}{1-\epsilon}}, \end{aligned} \quad (7)$$

where P_t^* is the price chosen by firms that can set their price optimally.⁷ Log-linearizing equation (7) around a zero-inflation steady state, the quasi-difference in aggregate inflation is related to the optimal relative price according to

$$\pi_t = \frac{\omega(1 - \theta)}{\theta + \omega(1 - \theta)} \pi_{t-1} + \frac{(1 - \omega)(1 - \theta)}{\theta + \omega(1 - \theta)} \widehat{p}_t^*, \quad (8)$$

where \widehat{p}_t^* denotes the percent deviation in p_t^* from $\bar{p}^* = 1$. Conditional on \widehat{p}_t^* , equation (8) implies that the correlation between inflation and its lag is an increasing function of ω and a decreasing function of θ . Further, conditional on lagged inflation, equation (8) implies that the correlation between π_t and \widehat{p}_t^* is a decreasing function of θ and ω .

⁶After all, not all price changes that show up in micro-data are necessarily determined optimally and some of the unchanged prices may in fact be optimal.

⁷Because real marginal costs are the same for all firms, in a symmetric equilibrium, firms that can set their price optimally will all choose the same price.

3.3 The pricing decision

I assume that $\omega + \theta > 0$, ruling out the case where all prices are flexible, but not ruling out the case where all firms change prices ($\theta = 0$) or the case where all price-changing firms optimize ($\omega = 0$). With this assumption, in period $t + 1$ a firm that cannot optimize its price between period t and period $t + 1$ will expect to charge the price

$$\begin{aligned} P_{t+1}(i) &= P_t(i) \left[\frac{\omega(1-\theta)}{\theta + \omega(1-\theta)} (1 + \pi_t) + \frac{\theta}{\theta + \omega(1-\theta)} \right], \\ &\equiv P_t(i) S_{t+1}, \end{aligned} \quad (9)$$

where the two terms in equation (9) correspond to the two non-optimizing pricing states, with each state weighted by its conditional probability. Iterating forward over equation (9), a firm that cannot optimally set its price will expect in period $t + j$ to charge the price

$$P_{t+j}(i) = P_t(i) \prod_{k=1}^j S_{t+k}. \quad (10)$$

Turning to the decision problem facing firms that can choose their price, in light of equation (10) these firms will choose $P_t(i)$ to maximize

$$\mathbb{E}_t \sum_{j=0}^{\infty} (\beta\mu)^j \frac{\lambda_{t+j}}{\lambda_t} Y_{t+j} \left[\left(\frac{P_t(i) \prod_{k=1}^j S_{t+k}}{P_{t+j}} \right)^{1-\epsilon} - mc_{t+j} \left(\frac{P_t(i) \prod_{k=1}^j S_{t+k}}{P_{t+j}} \right)^{-\epsilon} \right], \quad (11)$$

where $\mu \equiv \theta + \omega(1 - \theta)$ denotes the share of firms that cannot optimize their prices and λ_t is a shadow price representing the marginal utility of consumption in period t .

Differentiating equation (11) with respect to $P_t(i)$, the resulting first-order condition is

$$\mathbb{E}_t \sum_{j=0}^{\infty} (\beta\mu)^j \frac{\lambda_{t+j}}{\lambda_t} y_{t+j}(i) \left[\frac{p_t^*(i) \left(\prod_{k=1}^j S_{t+k} \right)}{\prod_{k=1}^j (1 + \pi_{t+k})} - \frac{\epsilon}{(\epsilon - 1)} mc_{t+j} \right] = 0, \quad (12)$$

which, when log-linearized around a zero-inflation steady state, yields

$$\widehat{p}_t^* = \beta\mu \mathbb{E}_t \widehat{p}_{t+1}^* + \beta\mu \left(\mathbb{E}_t \pi_{t+1} - \frac{\omega(1-\theta)}{\mu} \pi_t \right) + (1 - \beta\mu) \widehat{mc}_t. \quad (13)$$

Equation (13) establishes that, in addition to real marginal costs and its expected future price, the firm's pricing decision is shaped by current and expected future inflation. Because μ is increasing in both θ and ω , it is clear from equation (13) that increases in θ and ω raise the importance of future outcomes and lower the importance of current real marginal costs for the price chosen today.

3.4 The menu/information costs Phillips curve

To derive the Phillips curve, I combine equations (8) and (13) to obtain the expression

$$\begin{aligned} \pi_t = & \frac{\omega(1-\theta)}{\theta + \omega(1-\theta)(1+\beta)}\pi_{t-1} + \frac{\beta[\theta + \omega(1-\theta)]}{\theta + \omega(1-\theta)(1+\beta)}\mathbf{E}_t\pi_{t+1} \\ & + \frac{(1-\omega)(1-\theta)(1-\beta\mu)}{\theta + \omega(1-\theta)(1+\beta)}\widehat{mc}_t. \end{aligned} \quad (14)$$

Equation (14) has the form of a hybrid New Keynesian Phillips curve. Notice that the pricing parameters, θ and ω , affect both the lead-lag structure of inflation and the coefficient on real marginal costs. Specifically, it is not difficult to see that an increase in ω raises the coefficient on lagged inflation and lowers the coefficients on future inflation and real marginal costs. Similarly, an increase in θ raises the coefficient on future inflation and lowers the coefficients on lagged inflation and real marginal costs. Importantly then, a decline in the coefficient on real marginal costs need not imply higher menu costs (and greater price rigidity); it may, instead, imply higher information gathering/processing costs.

4 Some interesting special cases

It is interesting to relate the menu/information costs Phillips curve derived above to other specifications used in the literature. If I set $\omega = 0$, eliminating the pricing state in which firms index, then equation (14) collapses to

$$\pi_t = \beta\mathbf{E}_t\pi_{t+1} + \frac{(1-\theta)(1-\beta\theta)}{\theta}\widehat{mc}_t, \quad (15)$$

which is equivalent to the purely forward-looking new Keynesian Phillips curve associated with the Calvo (1983) model, equation (1). Similarly, if I set $\theta = 0$, eliminating the pricing state in which firms do not change prices, then equation (14) simplifies to

$$\pi_t = \frac{1}{1+\beta}\pi_{t-1} + \frac{\beta}{1+\beta}\mathbf{E}_t\pi_{t+1} + \frac{(1-\omega)(1-\beta\omega)}{\omega(1+\beta)}\widehat{mc}_t, \quad (16)$$

which is equivalent to the Christiano, Eichenbaum, and Evans (2005) full-indexation Phillips curve.

Since the menu/information costs Phillips curve encompasses both the Calvo (1983) model and the Christiano, Eichenbaum, and Evans (2005) model, it is natural to ask whether there might also be mathematical connections between equation (14) and the Smets and Wouters (2005) partial-indexation Phillips curve, which also encompasses these two specifications. To

address this question, note that, when approximated around a zero-inflation steady state, the partial-indexation Phillips curve is given by

$$\pi_t = \frac{\eta}{1 + \eta\beta} \pi_{t-1} + \frac{\beta}{1 + \eta\beta} \mathbb{E}_t \pi_{t+1} + \frac{(1 - \beta\xi)(1 - \xi)}{(1 + \eta\beta)\xi} \widehat{m\dot{c}}_t, \quad (17)$$

where $\eta \in [0, 1]$ represents the indexation parameter and $\xi \in (0, 1)$ represents the share of firms that can optimize their prices each period. As Smets and Wouters (2005) discuss, the model that underlies equation (17) is closely related to the Calvo (1983) model, with the modification that those firms that do not optimize their prices change their prices in proportion to lagged aggregate inflation.

Proposition 1: To a first-order (log-) approximation around a zero-inflation steady state, the partial-indexation Phillips curve and the menu/information costs Phillips curve are isomorphic.

Proof: Define $\eta \equiv \frac{\omega(1-\theta)}{\theta+\omega(1-\theta)}$ and $\xi \equiv \mu = \theta + \omega(1 - \theta)$, then the partial-indexation Phillips curve can be written as

$$\begin{aligned} \pi_t = & \frac{\frac{\omega(1-\theta)}{\theta+\omega(1-\theta)}}{1 + \beta \left(\frac{\omega(1-\theta)}{\theta+\omega(1-\theta)} \right)} \pi_{t-1} + \frac{\beta}{1 + \beta \left(\frac{\omega(1-\theta)}{\theta+\omega(1-\theta)} \right)} \mathbb{E}_t \pi_{t+1} \\ & + \frac{(1 - \omega)(1 - \theta)(1 - \beta\mu)}{\left[1 + \beta \left(\frac{\omega(1-\theta)}{\theta+\omega(1-\theta)} \right) \right] [\theta + \omega(1 - \theta)]} \widehat{m\dot{c}}_t. \end{aligned} \quad (18)$$

After some simple cancellations, equation (18) becomes

$$\begin{aligned} \pi_t = & \frac{\omega(1 - \theta)}{\theta + \omega(1 - \theta)(1 + \beta)} \pi_{t-1} + \frac{\beta}{\theta + \omega(1 - \theta)(1 + \beta)} \mathbb{E}_t \pi_{t+1} \\ & + \frac{(1 - \omega)(1 - \theta)(1 - \beta\mu)}{\theta + \omega(1 - \theta)(1 + \beta)} \widehat{m\dot{c}}_t, \end{aligned}$$

which has the same structure as the menu/information costs Phillips curve. Now, by inspection, for all $\omega \in [0, 1)$ and $\theta \in [0, 1)$ that satisfy $\omega + \theta > 0$, then $\eta \in [0, 1]$ and $\xi \in (0, 1)$, which establishes that the menu/information costs Phillips curve is a special case of the partial-indexation Phillips curve. Conversely, define $\theta \equiv \xi(1 - \eta)$ and $\omega \equiv \frac{\xi\eta}{1 - \xi(1 - \eta)}$, which imply $\mu = \xi$, then the menu/information costs Phillips curve can be written as

$$\begin{aligned} \pi_t = & \frac{\xi\eta}{\xi(1 - \eta) + \xi\eta(1 + \beta)} \pi_{t-1} + \frac{\beta[\xi(1 - \eta) + \xi\eta]}{\xi(1 - \eta) + \xi\eta(1 + \beta)} \mathbb{E}_t \pi_{t+1} \\ & + \frac{(1 - \xi)(1 - \beta\xi)}{\xi(1 - \eta) + \xi\eta(1 + \beta)} \widehat{m\dot{c}}_t, \end{aligned}$$

which in turn simplifies to

$$\pi_t = \frac{\eta}{1 + \eta\beta} \pi_{t-1} + \frac{\beta}{1 + \eta\beta} \mathbb{E}_t \pi_{t+1} + \frac{(1 - \xi)(1 - \beta\xi)}{(1 + \eta\beta)\xi} \widehat{mc}_t. \quad (19)$$

Equation (19) has the same structure as the partial-indexation Phillips curve. With respect to the parameter spaces, again by inspection, for all $\eta \in [0, 1]$ and $\xi \in (0, 1)$, then $\omega \in [0, 1)$ and $\theta \in [0, 1)$ and $\theta + \omega > 0$, which establishes that the partial-indexation Phillips curve is a special case of the menu/information costs Phillips curve. Since each specification is a special case of the other they must be isomorphic. ■

Proposition 1 establishes that the menu/information costs Phillips curve and the partial-indexation Phillips curve are mathematically equivalent, and this equivalence also has a strong intuition. The parameter η in the partial-indexation model has as its counterpart the convolution $\frac{\omega(1-\theta)}{\theta+\omega(1-\theta)}$ in the menu/information costs model. To appreciate why these two parameters play the same role, observe that the numerator of $\frac{\omega(1-\theta)}{\theta+\omega(1-\theta)}$ is the share of firms that index to lagged inflation and the denominator is the share of firms that are either indexing to lagged inflation or indexing to a zero inflation rate. In terms of the contribution to inflation being made by the non-optimizing firms, the convolution $\frac{\omega(1-\theta)}{\theta+\omega(1-\theta)}$ can be thought of as the weight on lagged inflation in a weighted average of lagged inflation and zero inflation, which is naturally equivalent to the weight on lagged inflation in a model with partial indexation. Similarly, it should be clear that the term $\frac{(1-\omega)(1-\theta)}{\theta+\omega(1-\theta)}$ regulates the relationship between inflation and the optimal relative price (see equation (8)) in the same way that $\frac{(1-\xi)}{\xi}$ does in the partial-indexation model and that these two expressions are equal when $\xi = \theta + \omega(1 - \theta)$, which is intuitive because ξ is the share of firms that do not optimize in the partial-indexation model and $\theta + \omega(1 - \theta)$ is the share of firms that do not optimize in the menu/information costs model.

5 System estimation

To estimate the menu/information costs Phillips curve I embed it within a small-scale dynamic stochastic general equilibrium (DSGE) model and estimate the resulting system using likelihood methods. The DSGE model is standard so I present the key equations without derivation. With c_t denoting consumption, R_t denoting the short-term nominal interest rate, and g_t denoting an aggregate consumption preference shock, the log-linearized consumption

Euler equation is given by (see Appendix C)

$$\widehat{c}_t = \frac{\gamma}{1+\gamma} \widehat{c}_{t-1} + \frac{1}{1+\gamma} \mathbf{E}_t \widehat{c}_{t+1} - \frac{(1-\gamma)}{\sigma} (R_t - \mathbf{E}_t \pi_{t+1} - \rho - g_t), \quad (20)$$

where $\gamma \in (0, 1)$ is the (external) habit parameter, $\sigma \in (0, \infty)$ is the coefficient of relative risk aversion, and $\rho = -\ln(\beta)$ is the discount rate. Combining the production technology, the resource constraint, and the household labor supply decision, real marginal costs are given by (see Appendix D)

$$\widehat{m}c_t = \left[\chi + \frac{\sigma}{(1-\gamma)} \right] \widehat{c}_t - \frac{\sigma\gamma}{(1-\gamma)} \widehat{c}_{t-1} - (1+\chi) u_t - g_t, \quad (21)$$

where $\chi \in (0, \infty)$ is the Frisch labor supply elasticity.

With respect to the nominal interest rate, I assume that R_t is set according to

$$R_t = (1 - \phi_r) [\rho + (1 - \phi_\pi) \bar{\pi} + \phi_\pi \mathbf{E}_t \pi_{t+1} + \phi_c \widehat{c}_t] + \phi_r R_{t-1} + \epsilon_t, \quad (22)$$

which is a standard forward-looking Taylor-type rule, essentially the same as the specification studied by Clarida, Galí, and Gertler (1998, 2000). Equation (22) postulates that the central bank responds with inertia to future expected inflation and, through consumption, to the state of the business cycle. Expected future inflation rather than current or lagged inflation enters the rule to capture the fact that central banks consider the future evolution of the economy when conducting monetary policy.

5.1 Maximum likelihood estimation

The model that I estimate has a rational expectations solution of the form

$$\mathbf{z}_t = \mathbf{h} + \mathbf{H}\mathbf{z}_{t-1} + \mathbf{G}\mathbf{v}_t, \quad (23)$$

where $\mathbf{z}_t = [\pi_t \quad \widehat{c}_t \quad R_t]'$ and $\mathbf{v}_t = [u_t \quad g_t \quad \epsilon_t]'$ and \mathbf{h} , \mathbf{H} , and \mathbf{G} are each functions of Γ , which denotes the vector of parameters to be estimated. By construction the eigenvalues of \mathbf{H} are bounded by 1 in modulus.

Because the rational expectations equilibrium for the model takes the form of equation (23), the assumption that $\mathbf{v}_t \sim n.i.i.d. [\mathbf{0}, \mathbf{\Omega}]$ implies that the concentrated log-likelihood function is

$$\log L_c(\Gamma; \{\mathbf{z}_t\}_2^T | \mathbf{z}_1) \propto (T-1) \ln \left[|(\mathbf{G})^{-1}| \right] - \frac{(T-1)}{2} \ln \left(|\widehat{\mathbf{\Omega}}| \right), \quad (24)$$

where

$$\hat{\Omega} = \sum_{t=2}^T \frac{[\mathbf{G}^{-1}(\mathbf{z}_t - \mathbf{h} - \mathbf{H}\mathbf{z}_{t-1})][\mathbf{G}^{-1}(\mathbf{z}_t - \mathbf{h} - \mathbf{H}\mathbf{z}_{t-1})]'}{T-1}.$$

Estimates of the parameters, $\mathbf{\Gamma}$, are found by maximizing equation (24), with their standard errors determined from the inverted Hessian evaluated at the optimum. The maximum likelihood estimates were obtained using a two-step approach. In the first step the genetic algorithm described in Duffy and McNelis (2001) was used. Briefly, this genetic algorithm is a stochastic search method that facilitates searching for a maximum over a wide parameter space. To implement the algorithm, a population of 1,000 initial candidate solutions was used, where the candidate solutions were drawn from a multivariate uniform distribution whose bounds were chosen to ensure that the model had a unique stable equilibrium within the search area. The genetic algorithm was allowed to run for a maximum of 300 generations or until each of the candidate solutions was identical to five decimal places, producing a set of first-step parameter estimates.⁸ In the second step, the parameter values that emerged from the genetic algorithm were used to initialize the BFGS optimization algorithm, which was iterated to convergence. The maximum likelihood estimates that I report, and their standard errors, reflect the maximum obtained by the BFGS algorithm. This two-step approach allows me to search widely over the likelihood function, helping ensure that a global maximum rather than a local maximum is located.

5.2 Bayesian estimation

Let \mathbf{M} denote the model space and $m_j \in \mathbf{M}$, $j \in \{1, 2, \dots, M\}$, reference an arbitrary model. With the parameters of model m_j represented by $\mathbf{\Gamma}_{m_j}$, $p(\mathbf{\Gamma}_{m_j} | m_j)$ is the prior density for $\mathbf{\Gamma}_{m_j}$, $p(\{\mathbf{z}_{1t}\}_2^T | \mathbf{\Gamma}_{m_j}, m_j)$ is the conditional data density, and $p(\mathbf{\Gamma}_{m_j} | \{\mathbf{z}_{1t}\}_2^T, m_j)$ is the posterior density of the parameter density conditional on the data and the model. As always with Bayesian estimation, interest centers on the posterior density, which from Bayes's theorem, is given by

$$p(\mathbf{\Gamma}_{m_j} | \{\mathbf{z}_t\}_2^T, m_j) = \frac{p(\{\mathbf{z}_t\}_2^T | \mathbf{\Gamma}_{m_j}, m_j) p(\mathbf{\Gamma}_{m_j} | m_j)}{p(\{\mathbf{z}_t\}_2^T | m_j)}. \quad (25)$$

⁸Advantages to using a genetic algorithm are that it does not require taking numerical derivatives and that, by sampling over the entire admissible parameter space, it helps to ensure that a global maximum of the likelihood function is obtained. Other than the fact that "mutation" was not applied in the creation of "children," the genetic algorithm employed in this paper follows precisely that described in Appendix A of Duffy and McNelis (2001), to which interested readers are referred. Mutation was not applied since the solution obtained from the genetic algorithm was not the final estimate, but rather only the source of starting values for a quasi-Newton hill climber (BFGS).

To draw from the posterior density, I use the random walk chain Metropolis-Hastings algorithm. Ten overdispersed chains of length 60,000 were constructed from which the first 10,000 “burn-in” draws were discarded, leaving a total of 500,000 usable draws. Convergence of the chains was determined using diagnostics developed by Gelman (1995) and Geweke (1992).

For the model comparison exercise in Section 6, I use Geweke’s (1999) modification of the Gelfand and Dey (1994) method to calculate the marginal data density, or marginal likelihood,

$$p\left(\{z_t\}_2^T | m_j\right) = \int_{\mathbf{\Gamma}_{m_j}} p\left(\{z_t\}_2^T | \mathbf{\Gamma}_{m_j}, m_j\right) p\left(\mathbf{\Gamma}_{m_j} | m_j\right) d\mathbf{\Gamma}_{m_j}, \quad (26)$$

which is the probability of observing the data given model m_j . As equation (26) shows, the marginal likelihood is evaluated by averaging the conditional data density with respect to the prior density. After evaluating the marginal likelihood for each model, the posterior probability associated with model $m_k \in \mathbf{M}$ is calculated according to

$$p\left(m_k | \{z_t\}_2^T\right) = \frac{p\left(\{z_t\}_2^T | m_k\right) p\left(m_k\right)}{\sum_{j=1}^M p\left(\{z_t\}_2^T | m_j\right) p\left(m_j\right)}, \quad (27)$$

where $p(m_j)$ is the prior probability associated with model m_j , $j \in \{1, 2, \dots, M\}$. With respect to the model space, \mathbf{M} , I apply a discrete uniform prior to the model space, thus $p(m_j) = \frac{1}{M}$, $j \in \{1, \dots, M\}$.

5.3 Data

To estimate the model, I use U.S. data spanning the period 1982:1 – 2002:4, which excludes the period of nonborrowed reserves targeting that occurred in the early 1980s, but otherwise reflects the time during which Paul Volcker and Alan Greenspan were Federal Reserve chairmen. I use the quarterly average of the federal funds rate to measure R_t , use $100 \times \ln(C_t/C_t^T)$ to measure the consumption gap, where C_t is real consumption and C_t^T is trend consumption,⁹ and use $400 \times \ln(P_t/P_{t-1})$, where P_t is the personal consumption expenditure (PCE) price index, to measure inflation.

5.4 Priors

Aside from the parameters describing the shock processes, the key model parameters are $\mathbf{\Gamma} = \{\chi, \rho, \gamma, \sigma, \theta, \omega, \bar{\pi}, \phi_\pi, \phi_c, \phi_r\}$. However, the data are sufficiently uninformative of the labor

⁹Trend consumption was constructed using the Hodrick-Prescott filter with $\lambda = 1,600$.

supply elasticity, χ , that precise estimates could not be obtained using maximum likelihood. As a consequence, and to enable comparison between the FIML and the Bayesian estimates, I set χ equal to 2.50 during estimation, with this value based on Smets and Wouters (2005) and Chang and Kim (2005). The priors for the remaining behavioral parameters are summarized in Table 2a.

Parameter	Distribution	Mean	Std Dev.	90% Interval
ρ	Normal	2.50	0.50	[1.68, 3.32]
γ	Beta	0.75	0.10	[0.57, 0.90]
σ	Gamma	2.00	2.00	[0.10, 5.99]
θ	Beta	0.66	0.10	[0.49, 0.82]
ω	Beta	0.50	0.20	[0.25, 0.75]
$\bar{\pi}$	Normal	3.00	0.50	[2.18, 3.82]
ϕ_{π}	Normal	1.50	0.20	[1.17, 1.83]
ϕ_c	Normal	1.00	0.20	[0.67, 1.33]
ϕ_r	Beta	0.75	0.10	[0.57, 0.90]

Briefly, the priors for ρ and $\bar{\pi}$ have means equaling 2.50 and 3.00 percent, respectively, at annual rates. The priors for γ and ϕ_r are each Beta distributions with means equaling 0.75. The prior for the inflation indexation parameter, ω , is a uniform distribution over the unit interval while that for the frequency of price adjustment, θ , has a Beta distribution that is centered on 0.66, reflecting the results in Nakamura and Steinsson (2006). The prior for the coefficient of relative risk aversion, σ , has a Gamma distribution with a mean equaling 2.00 and, to reflect the wide range of estimates in the literature, a relatively large standard deviation of 2.00.

The prior for the shock process was implemented as follows. First, the solution to the rational expectations model, equation (23), was written as

$$\mathbf{z}_t = \mathbf{h} + \mathbf{H}\mathbf{z}_{t-1} + \boldsymbol{\varepsilon}_t, \quad (28)$$

where $\boldsymbol{\varepsilon}_t = \mathbf{G}\mathbf{v}_t$ are reduced-form (supply, demand, and policy) shocks. The priors for the elements in $\boldsymbol{\Sigma} = \mathbf{E}[\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t']$ are summarized in Table 2b.

Parameter	Distribution	Mean	Std Dev.	90% Interval
σ_1	Inverted Gamma	1.00	0.20	[0.72, 1.36]
σ_2	Inverted Gamma	0.50	0.20	[0.27, 0.87]
σ_3	Inverted Gamma	0.70	0.20	[0.44, 1.07]
$\text{cov}(\varepsilon_1\varepsilon_2)$	Normal	0.00	0.20	[-0.33, 0.33]
$\text{cov}(\varepsilon_1\varepsilon_3)$	Normal	0.00	0.20	[-0.33, 0.33]
$\text{cov}(\varepsilon_2\varepsilon_3)$	Normal	0.00	0.20	[-0.33, 0.33]

6 Estimates of the menu/information costs model

Table 3 presents my estimates of the DSGE model parameters. The table displays FIML estimates, with standard errors in parentheses,¹⁰ along with the posterior mean, median, and mode, and a 90 percent probability interval. Also shown are the maximized value of the (log-) likelihood function (log-L) and the (log-) marginal likelihood (log-ML). The FIML estimates are shown so that interested readers can easily assess the prior's role in shaping the Bayesian estimates.

Parameter	FIML	Post. Mean	Post. Median	Post. Mode	90% Interval
ρ	2.489 (0.741)	2.566	2.573	2.569	[1.932, 3.170]
γ	0.820 (0.083)	0.820	0.822	0.838	[0.731, 0.903]
σ	5.860 (5.624)	4.471	4.083	2.969	[1.698, 8.609]
θ	0.305 (0.073)	0.407	0.405	0.407	[0.308, 0.510]
ω	0.971 (0.014)	0.950	0.952	0.953	[0.926, 0.970]
$\bar{\pi}$	3.339 (1.402)	3.287	3.278	3.338	[2.782, 3.822]
ϕ_{π}	1.625 (1.164)	1.591	1.584	1.580	[1.318, 1.886]
ϕ_c	1.222 (1.174)	1.000	1.000	0.997	[0.680, 1.321]
ϕ_r	0.868 (0.039)	0.846	0.847	0.849	[0.802, 0.886]
log- L	-262.388	-			
log- ML	-	-289.795			

Turning to the coefficient values themselves, the FIML and Bayesian estimates of the rate of time preference, ρ , are about 2.5. These values are consistent with estimates of the equilibrium real interest rate (Laubach and Williams, 2003) and place the quarterly discount factor at just over 0.99, in line with values widely used in calibration exercises.

Looking at the utility function parameters, the habit formation parameter, γ , is estimated to be about 0.82, implying that habit formation is important and that there is considerable inertia in consumption. Elsewhere, estimates of γ vary widely. Smets and Wouters (2005) estimate $\gamma = 0.69$, Altig, Christiano, Eichenbaum, and Linde (2004) estimate $\gamma = 0.65$, Giannoni and Woodford (2003) estimate $\gamma = 1.00$, while the results in Smets (2003) and Cho and Moreno (2005) imply that γ equals 0.79 and 1.00, respectively. Calibration exercises

¹⁰These standard errors were calculated from the inverted Hessian matrix evaluated at the maximum of the likelihood function.

often set γ to 0.80 (McCallum and Nelson, 1999). Turning to σ , estimates in the literature are also wide-ranging. The FIML estimate is imprecise but places σ at about 5.86, while the Bayesian estimation returns a posterior median for σ that is about 4.1. Elsewhere, Fuhrer (2000) obtains $\sigma = 6.11$ while Kim (2000) obtains $\sigma = 14.22$. Using Bayesian methods, Smets and Wouters (2005) get $\sigma = 1.62$ for the posterior median, while Levin, Onatski, Williams, and Williams (2006) report $\sigma = 2.19$ for the posterior mean. At the other end of the spectrum, Rotemberg and Woodford (1997) estimate $\sigma = 0.16$, Amato and Laubach (2003) estimate $\sigma = 0.26$, and Giannoni and Woodford (2003) estimate $\sigma = 0.75$.

Regarding the policy-rule parameters, the FIML and Bayesian estimates place the implicit inflation target at around 3.3 percent. These estimates of $\bar{\pi}$ are very similar to those obtained by Clarida, Galí, and Gertler (2000), while being slightly higher than that obtained by Favero and Rovelli (2003), who estimate $\bar{\pi}$ to be 2.63 percent. I estimate the coefficient on expected future inflation to be about 1.6, the coefficient on the consumption gap to be about 1, and the coefficient on lagged interest rates to be about 0.85. These coefficients are all consistent with other estimated Taylor-type rules (see Clarida, Galí, and Gertler (2000) and Dennis (2006a)); they indicate an activist, but inertial, approach to monetary policy and rule out sunspot behavior.

With respect to pricing behavior, the two key parameters are θ and ω . The FIML estimate of θ is 0.305, while a two-standard-deviation confidence interval spans 0.159 to 0.451. The Bayesian estimation has the distribution for θ centered on about 0.407, with a 90 percent probability interval covering 0.308 to 0.510. These estimates place the frequency of price adjustment somewhere around 0.60, representing relatively frequent price adjustment. Clearly, one of the main findings that emerges from the estimates in Table 3 is that movements in the PCE price index are consistent with frequent price adjustment.

Finally, because the estimates of θ reveal that firms do, in fact, change prices quite frequently, they suggest that menu costs are not a huge impediment to a firm changing its price. At the same time, the estimates of ω are large, and they imply that most firms that change prices do so using indexation, a result that is consistent with the view that information gathering/processing costs are more important for pricing behavior than menu costs.

6.1 The partial-indexation model and the frequency of price adjustment

Proposition 1 shows how estimates of the partial-indexation Phillips curve can be used to shed light on the values of θ and ω in the menu/information costs Phillips curve; a range of

estimates for the U.S are reported in Table 4.¹¹

Study	Sample	ξ	η	θ	ω
Sahuc (2004)	1970:1 – 2002:4	0.93	0.67	0.31	0.90
Smets & Wouters (2005)	1974:1 – 2002:2	0.87	0.66	0.30	0.82
Rabanal & Rubio-Ramírez (2005)	1960:1 – 2001:4	0.84	0.76	0.23	0.96
Dennis (2006b)	1983:1 – 2004:2	0.88	0.69	0.27	0.83

The key result that emerges from Table 4 is that estimates of the partial-indexation Phillips curve are consistent with those for the menu/information costs model shown in Table 3, even though they are obtained from DSGE models that are quite different from the one I estimate. Specifically, the estimates of θ in Table 4 imply that the frequency of price adjustment is around 0.23 – 0.31, consistent with quite rapid price adjustment; the estimates of ω indicate that the share of firms that index prices is large, suggesting that price indexation explains the majority of price changes.

6.2 Alternative pricing specifications

In this section I estimate four additional Phillips curve specifications and consider the results in light of the menu/information costs model estimated above. With the remainder of the model continuing to be given by equations (20), (21), and (22), I estimate the New Keynesian Phillips curve, equation (15), the full-indexation Phillips curve, equation (16), the partial-indexation Phillips curve, equation (17), and a hybrid Phillips curve developed by Galí and Gertler (1999), which is given by

$$\pi_t = \frac{\omega}{\theta + \omega [1 - \theta (1 - \beta)]} \pi_{t-1} + \frac{\beta\theta}{\theta + \omega [1 - \theta (1 - \beta)]} \mathbf{E}_t \pi_{t+1} + \frac{(1 - \omega)(1 - \theta)(1 - \beta\theta)}{\theta + \omega [1 - \theta (1 - \beta)]} \widehat{mc}_t. \quad (29)$$

Although θ and ω carry the same interpretation in the Galí and Gertler (1999) model as they do in the menu/information costs model, the firm’s optimization problem is quite different. Thus, while similar in the variables it depends upon, equation (29) is otherwise quite dissimilar to the menu/information costs Phillips curve, equation (14).

¹¹As earlier, where necessary, the estimates in Table 4 have been made consistent with a Cobb-Douglas production technology and rental markets for capital and labor.

6.2.1 Estimates

Table 5 presents median estimates of the parameters in the four alternative pricing specifications discussed above.¹² As noted previously, these alternative pricing specifications are estimated jointly with equations (20), (21), and (22).

Parameter	Calvo	Full Indexation.	Galí-Gertler	Partial Indexation.
ρ	2.521	2.659	2.702	2.645
γ	0.802	0.841	0.839	0.832
σ	3.657	4.515	3.956	4.048
θ	0.959	—	0.905	0.282
ω	—	0.944	0.757	0.959
$\bar{\pi}$	3.472	2.985	3.048	3.130
ϕ_π	1.597	1.531	1.541	1.546
ϕ_c	1.001	1.011	1.023	1.014
ϕ_r	0.835	0.836	0.833	0.841
log- ML	-325.467	-289.204	-288.094	-287.287

Setting the estimates of θ and ω aside for the moment, the estimates of the remaining structural parameters are both reasonably consistent across the three specifications and with those for the menu/information costs model shown in Table 3. The estimates of the habit formation parameter, γ and the coefficient of relative risk aversion, σ are about 0.8 and 4.0, respectively, in line with those obtained earlier. Further, although there is some variation in the estimates of the inflation target, the estimates place it around 3.0 percent, and the estimates of the remaining policy rule parameters are all in line with those shown in Table 3.

With respect to the two pricing parameters, estimates from the Calvo model, the full-indexation model, and the Galí-Gertler model all suggest that the share of firms that set their price optimally each quarter is small. However, where the Calvo model suggests that the majority of firms keeps their prices unchanged, the full-indexation model suggests that the same majority index their prices to lagged inflation. The Galí-Gertler model suggests that while the majority of firms keeps their prices unchanged each period, about 76 percent of those that do change their prices do so by rule of thumb.

Although the Galí-Gertler model and the menu/information costs model both highlight the importance of rule-of-thumb pricing or price indexation for inflation dynamics, they tell

¹²For the model containing the partial-indexation Phillips curve, the prior for ξ and η was a Beta distribution with mean of 0.5 and standard deviation of 0.2. The prior for the remaining parameters was the same as that reported in Tables 2a and 2b.

very different stories regarding the frequency of price adjustment. Specifically, the Galí-Gertler model indicates that about 90 percent of firms keep their price unchanged each quarter whereas the menu/information costs model places this share at about 40 percent. However, the estimates of θ from the Galí-Gertler model are quite counterintuitive. To see why, consider equation (30),

$$\pi_t = 0.457\pi_{t-1} + 0.542E_t\pi_{t+1} + 0.001\widehat{m}c_t, \quad (30)$$

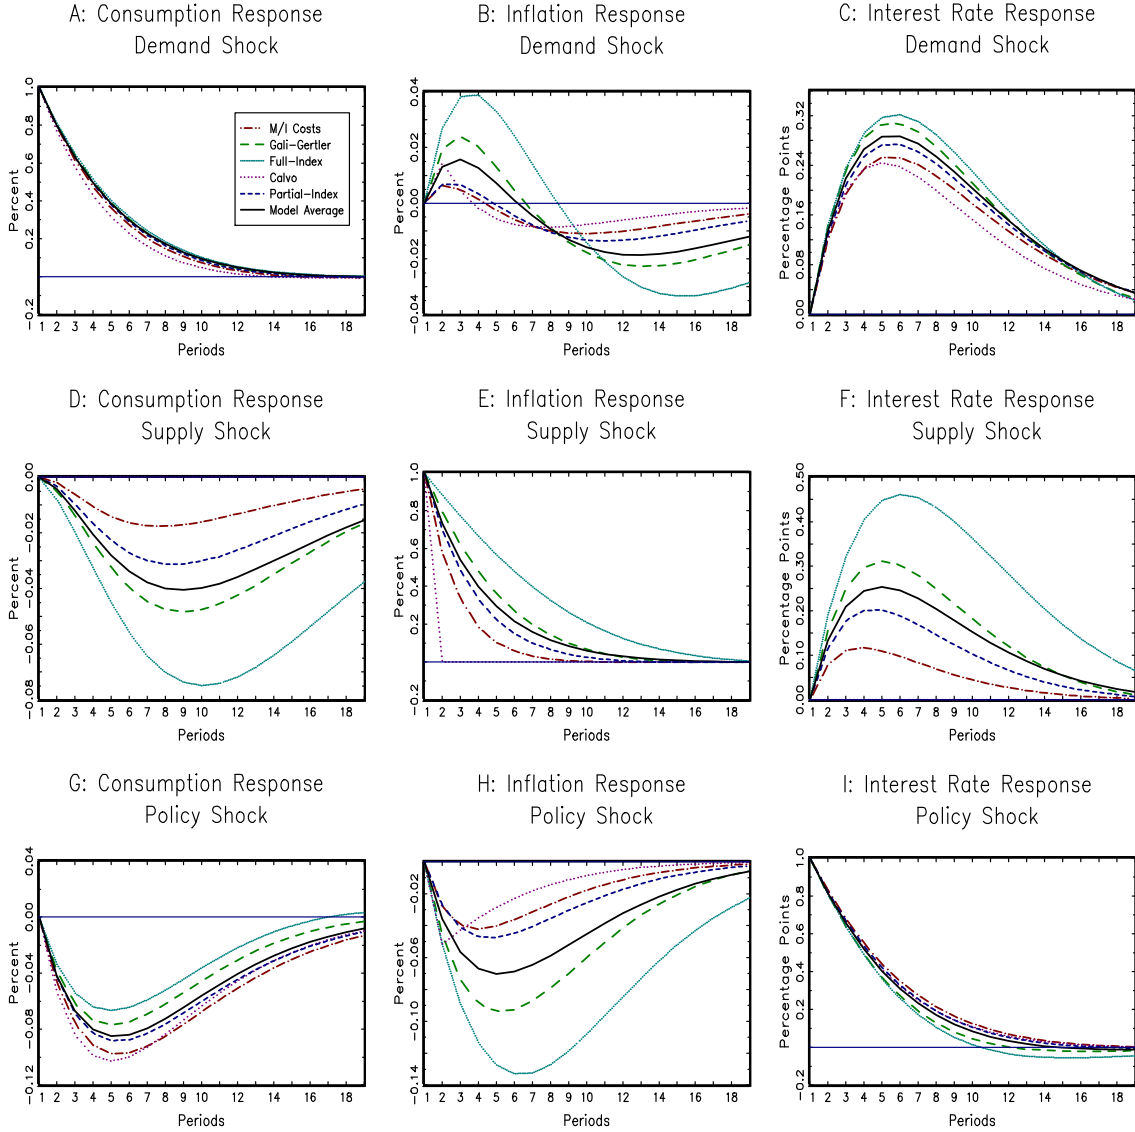
which reports the empirical representation of the Galí-Gertler Phillips curve. According to the Galí-Gertler Phillips curve, it is the behavior of just 2 percent of firms that accounts for all of the endogenous persistence in inflation and, hence, for the coefficient of 0.46 on lagged inflation. By way of contrast, the empirical representation of the menu/information costs model is

$$\pi_t = 0.369\pi_{t-1} + 0.629E_t\pi_{t+1} + 0.001\widehat{m}c_t, \quad (31)$$

and here the coefficient of 0.37 on lagged inflation is important because about 57 percent of firms change prices using indexation.

7 Business cycle dynamics

In this section I demonstrate that the differences between the four pricing models estimated above are not just statistical, nor are they just theoretical, rather they are economically important. I study how the models respond to shocks, considering demand shocks, supply shocks, and monetary policy shocks. Furthermore, I use the Bayesian estimates to construct predictive densities for each model and for each shock, and, exploiting the posterior model probabilities, employ Bayesian model averaging to examine the (weighted) average response to each shock. The results for unit shocks are shown in Figure 1, which plots the median of the predictive densities for each model, together with the results from the Bayesian model averaging exercise.



Predictive densities following unit shocks

To understand these shock responses it is useful to first focus on the Bayesian model averaging exercise. Panels A to C correspond to the demand shock, Panels D to F correspond to the technology shock, and Panels G to I correspond to the monetary policy shock. Following a positive demand shock, households take advantage of the fact that higher utility can be achieved by consuming more now and increase their labor supply in order to raise their income to facilitate greater consumption expenditures (Panel A). The increase in labor supply is offset by a rise in labor demand, as firms increase production to meet rising demand, and, on balance, the market-clearing real wage increases, raising real marginal costs and causing a small and

gradual increase in aggregate inflation (Panel B). Faced with stronger consumer demand, and slightly higher inflation, the central bank raises interest rates (Panel C).

Following a positive supply shock, the marginal product of labor increases, pushing up the demand for labor and raising the market-clearing real wage. Household income rises due to the higher real wage and from households increasing their hours worked, which pushes up consumption (Panel D). At the same time, the increase in the marginal product of labor has the effect of lowering real marginal costs, which puts downward pressure on inflation (Panel E). The decline in inflation is substantial and the policy response is to lower nominal, and hence also real, interest rates (Panel F). Finally, following a positive monetary policy shock, real interest rates rise (Panel I), which induces households to defer consumption. To offset the fall in consumption (Panel G), households increase their labor supply, which puts downward pressure on real wages. Although firms respond to declining demand by cutting their demand for labor, the market-clearing real wage falls, lowering real marginal costs and inflation (Panel H).

Relative to the menu/information costs model, the poor performance of the Calvo model is clear. Following a supply shock (Panel E) inflation falls, but then immediately returns to baseline, without any effect on consumption or interest rates. More generally, the Calvo model's responses to all three shocks differ importantly from the responses of the menu/information costs model in that they are not "hump-shaped," underscoring the Estrella and Fuhrer (2002) criticism of the New Keynesian Phillips curve. With regard to the full-indexation model, although its responses are hump-shaped, they are also generally much larger than those of either the menu/information costs model or the Bayesian model average. These large responses are particularly evident in how the model behaves following a supply shock (Panels D to F), but are also evident in Panels B and H. The source of these large responses is the fact that many firms index their prices to lagged inflation and no firms keep their prices unchanged following shocks. The large inflation responses give rise to large interest rate responses, which, in turn, generate relatively large consumption responses by households. It is clear from the posterior model probabilities and the behavior of the Bayesian model average that the data provide considerably less support for the behavior of the full-indexation model than they do for the behavior of the menu/information costs model or the partial-indexation model. Lastly, it is important to note that the menu/information costs model and the Galí-Gertler model behave similarly, both qualitatively and quantitatively.

8 Conclusion

The Calvo-based New Keynesian Phillips curve has been widely criticized for being economically implausible, for being inconsistent with micro-data on the frequency of price adjustment, and for being unable to account for the persistence in inflation. Popular alternatives, such as the full-indexation model and the partial-indexation model, are much better able to explain the persistence in inflation, but, because they assume that all prices change every period, they too are economically implausible and are unable to match micro-evidence on the frequency of price changes. These criticisms are important because New Keynesian business cycle models are increasingly used to study issues such as how monetary policy should be conducted to maximize welfare, and the nature of these policies hinge critically on precisely how and why prices are rigid. More generally, they challenge whether the leading New Keynesian models of price adjustment provide a useful and economically sensible description of inflation dynamics. Against this background, the main contribution of this paper is to present a pricing model that can usefully speak to these criticisms.

I begin by presenting estimates of the frequency of price adjustment obtained when the Calvo-based New Keynesian Phillips curve is estimated in isolation, outlining the implications these estimates have for the average duration between price changes. Next, I introduce a pricing model in which each period a share of firms get to change their prices and within this share a proportion changes their prices optimally while the remaining proportion changes their prices using an indexation rule. I highlight that when it comes to reconciling macro- and micro-evidence on the frequency of price adjustment and to accounting for the persistence in inflation, the model I develop holds many attractions. First, in this model not all prices necessarily change every period and when prices change they do not necessarily change optimally. Second, the model's share parameters can be interpreted easily in light of the costs firms face when changing prices. Traditional menu cost factors, which affect all price changes not just optimal price changes, are readily associated with the share of firms that change their prices. Similarly, information gathering/processing costs are readily associated with the share of price-changing firms that resort to price indexation. Third, because some firms index their prices to lagged inflation, the model has a mechanism for generating intrinsic inflation persistence.

After outlining the menu/information costs pricing model, I derive its Phillips curve and relate it to other specifications in the literature. Specifically, I prove that it encompasses the Calvo Phillips curve, the full-indexation Phillips curve, and the partial-indexation Phillips

curve. This encompassing result, together with the fact that the full- and the partial-indexation models counterfactually force all firms to change their prices every period, makes the menu/information costs Phillips curve particularly attractive for empirical applications. Taking this as motivation, I build a small-scale New Keynesian business cycle model and estimate versions of it on U.S. macroeconomic data.

My main empirical results are as follows. First, my estimates of the menu/information costs model place the quarterly frequency of price adjustment at about 0.6. In this respect, the menu/information costs model, unlike many other time-contingent pricing models, reveals a relatively high frequency of price adjustment. Second, with around 60 percent of firms changing their prices each quarter and with over 90 percent of them resorting to indexation, my estimates are consistent with the widely held view that menu costs are a much less important factor for price setting than information gathering/processing costs. Importantly, the estimates of the frequency of price adjustment and the share of indexing firms that I obtain are supported by estimates of the partial-indexation Phillips curve, are robust to whether the model is estimated using FIML or Bayesian methods, and are consistent with relatively frequent price adjustment.

A Appendix: Aggregate prices

From the Dixit-Stiglitz aggregator, the price level is defined according to

$$P_t \equiv \left[\int_0^1 P_t(i)^{1-\varepsilon} di \right]^{\frac{1}{1-\varepsilon}}. \quad (\text{A1})$$

Recognizing that at any point in time firms either set their price optimally, use price indexation, or keep their price unchanged, equation (A1) is equivalent to

$$P_t^{1-\varepsilon} = (1-\omega)(1-\theta)P_t^{*1-\varepsilon} + (1+\pi_{t-1})^{1-\varepsilon} \int_0^{\theta+\omega(1-\theta)} P_{t-1}(i)^{1-\varepsilon} di + \int_0^\theta P_{t-1}(i)^{1-\varepsilon} di. \quad (\text{A2})$$

Because the firms that do not change their prices and that use indexation are chosen randomly, equation (A2) is equivalent to

$$\begin{aligned} P_t^{1-\varepsilon} &= (1-\omega)(1-\theta)P_t^{*1-\varepsilon} + (1+\pi_{t-1})^{1-\varepsilon} \omega(1-\theta) \int_0^1 P_{t-1}(i)^{1-\varepsilon} di \\ &\quad + \theta \int_0^1 P_{t-1}(i)^{1-\varepsilon} di, \\ &= (1-\omega)(1-\theta)P_t^{*1-\varepsilon} + \omega(1-\theta)(1+\pi_{t-1})^{1-\varepsilon} P_{t-1}^{1-\varepsilon} + \theta P_{t-1}^{1-\varepsilon}. \end{aligned} \quad (\text{A3})$$

Finally, equation (A3) implies

$$P_t = \left[(1 - \omega)(1 - \theta) P_t^{*1-\varepsilon} + \omega(1 - \theta)(1 + \pi_{t-1})^{1-\varepsilon} P_{t-1}^{1-\varepsilon} + \theta P_{t-1}^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}}, \quad (\text{A4})$$

which is equation (7) in the text. If I assume, say, that there exists an initial period in which there is no relative price dispersion, then it is straightforward to see that equation (A4) holds for all $t \geq 1$.

Alternatively, by accounting for how firms that either did not change prices today or that used price indexation today have set prices in the past, and exploiting, first, that a firm's pricing strategy is determined randomly, and, second, that as the initial period tends to $-\infty$ the share of firms setting prices today that have *never* set their prices optimally converges to zero (or that there exists an initial period in which there is no relative price dispersion, making the summation in equation (A5) finite), equation (A2) can be written as

$$P_t^{1-\varepsilon} = (1 - \omega)(1 - \theta) \sum_{k=0}^{\infty} \left(\prod_{l=1}^k [\theta + \omega(1 - \theta)(1 + \pi_{t+l-1})] \right) P_t^{*1-\varepsilon}, \quad (\text{A5})$$

which, under these assumptions, is convergent for all $\omega \in [0, 1)$ and $\theta \in [0, 1)$. From equation (A5), it follows that $P_{t-1}^{1-\varepsilon}$ is given by

$$P_{t-1}^{1-\varepsilon} = (1 - \omega)(1 - \theta) \sum_{k=0}^{\infty} \left(\prod_{l=1}^k [\theta + \omega(1 - \theta)(1 + \pi_{t+l-1})^{1-\varepsilon}] \right) P_{t-1}^{*1-\varepsilon}, \quad (\text{A6})$$

and combining equations (A5) and (A6) yields

$$P_t^{1-\varepsilon} - [\theta + \omega(1 - \theta)(1 + \pi_{t+l-1})^{1-\varepsilon}] P_{t-1}^{1-\varepsilon} = (1 - \omega)(1 - \theta) P_t^{*1-\varepsilon},$$

which implies

$$\begin{aligned} P_t^{1-\varepsilon} &= (1 - \omega)(1 - \theta) P_t^{*1-\varepsilon} + [\theta + \omega(1 - \theta)(1 + \pi_{t+l-1})^{1-\varepsilon}] P_{t-1}^{1-\varepsilon}, \\ &= (1 - \omega)(1 - \theta) P_t^{*1-\varepsilon} + \omega(1 - \theta)(1 + \pi_{t+l-1})^{1-\varepsilon} P_{t-1}^{1-\varepsilon} + \theta P_{t-1}^{1-\varepsilon}, \end{aligned}$$

leading to

$$P_t = \left[(1 - \omega)(1 - \theta) P_t^{*1-\varepsilon} + \omega(1 - \theta)(1 + \pi_{t+l-1})^{1-\varepsilon} P_{t-1}^{1-\varepsilon} + \theta P_{t-1}^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}},$$

which, again, is equation (7) in the text.

B Appendix: The non-zero-inflation steady state case

Let $\bar{\pi} \in (-1, \tilde{\pi}]$ denote the inflation rate in the nonstochastic steady state. With $p_t^* = \frac{P_t^*}{P_t}$ representing the optimal relative price, it follows from equation (A4) (equation (7) in the text) that in a non-zero-inflation steady state, the steady state optimal relative price, \bar{p}^* , is given by

$$\bar{p}^* = \left[\frac{1 - \omega(1 - \theta) - \theta(1 + \bar{\pi})^{\varepsilon-1}}{(1 - \omega)(1 - \theta)} \right]^{\frac{1}{1-\varepsilon}}.$$

In order for the steady state optimal relative price to be positive ($\bar{p}^* > 0$), it must be the case that

$$1 - \omega(1 - \theta) - \theta(1 + \bar{\pi})^{\varepsilon-1} > 0, \quad (\text{B1})$$

which, for $\theta \in (0, 1)$, leads to

$$\tilde{\pi}_1 < \left[\left(\frac{1 - \omega(1 - \theta)}{\theta} \right)^{\left(\frac{1}{\varepsilon - 1} \right)} - 1 \right]. \quad (\text{B2})$$

Equation (B1) implies that \bar{p}^* is increasing in $\bar{\pi}$ for $\bar{\pi} \in (-1, \tilde{\pi}_1]$, and greater than one for $\bar{\pi} \in (0, \tilde{\pi}_1]$, less than one for $\bar{\pi} \in (-1, 0)$, and equal to one for $\bar{\pi} = 0$. Notably, because prices are not fully flexible, on occasions when they can change prices, firms respond to a positive steady state inflation rate by raising prices by more than they otherwise would have, with inflation then eroding these high relative prices over time. Of course, the extent to which a positive steady state inflation rate lifts the steady state optimal relative price is mitigated by greater substitutability between goods and by greater price flexibility. Now, looking at equation (B2), it is straightforward to see that $\tilde{\pi}_1$ is decreasing in ε , which is intuitive because greater substitutability between goods is incompatible with sticky prices unless steady state inflation is low. Similarly, for a given elasticity of substitution between goods, greater price rigidity requires a lower steady state inflation rate if the optimal relative price is to remain well-defined.

Log-linearizing equation (A4) around a non-zero-inflation steady state, the quasi-difference in aggregate inflation is related to the optimal relative price according to

$$\begin{aligned} \pi_t - \bar{\pi} &= \frac{\omega(1 - \theta)}{\omega(1 - \theta) + \theta(1 + \bar{\pi})^{\varepsilon - 1}} (\pi_{t-1} - \bar{\pi}) \\ &+ (1 + \bar{\pi}) \left(\frac{1 - \omega(1 - \theta) - \theta(1 + \bar{\pi})^{\varepsilon - 1}}{\omega(1 - \theta) + \theta(1 + \bar{\pi})^{\varepsilon - 1}} \right) \hat{p}_t^*, \end{aligned} \quad (\text{B3})$$

where \hat{p}_t^* denotes the percent deviation in p_t^* from \bar{p}^* . Equation (B3) is the analogue of equation (8) in the text.

From equation (11), the pricing decision for firms that can set their price optimally leads to the first-order condition

$$E_t \sum_{j=0}^{\infty} (\beta\mu)^j \frac{\lambda_{t+j}}{\lambda_t} y_{t+j}(i) \left[\frac{p_t^*(i) \left(\prod_{k=1}^j S_{t+k} \right)}{\prod_{k=1}^j (1 + \pi_{t+k})} - \frac{\epsilon}{(\epsilon - 1)} mc_{t+j} \right] = 0, \quad (\text{B4})$$

which, provided $\beta\mu(1 + \bar{\pi}) < 1$, implies that the steady state relationship between real marginal costs and the optimal relative price is

$$\bar{m}\bar{c} = \frac{\varepsilon - 1}{\varepsilon} \left(\frac{\omega(1 - \theta)}{\theta + \omega(1 - \theta)} + \frac{\theta}{\theta + \omega(1 - \theta)} \frac{1 - \beta\mu(1 + \bar{\pi})}{1 - \beta\mu} \right) \bar{p}^*. \quad (\text{B5})$$

Equation (B5) reveals the steady state markup over real marginal costs as a function of the model's parameters. In a zero-inflation steady state, the markup depends only on the elasticity of substitution between goods and is given by $\frac{1}{\varepsilon - 1}$. More generally, equation (B5) shows that in a non-zero-inflation steady state the markup depends on β , θ , ω , and $\bar{\pi}$, in addition to ε . Although the steady state markup continues to be an increasing function of ε , it is also an increasing function of the steady state inflation rate.

The restriction that the parameters satisfy $\beta\mu(1 + \bar{\pi}) < 1$, which is needed to ensure that the model has a well-defined steady state, gives rise to the condition

$$\tilde{\pi}_2 < \left(\frac{1}{\beta\mu} \right) - 1, \quad (\text{B6})$$

indicating that $\tilde{\pi}_2$ is declining in the discount factor, β , and in the share of non-optimizing firms, μ . Combining equations (B2) and (B6), $\tilde{\pi}$ is given by

$$\tilde{\pi} = \min\{\tilde{\pi}_1, \tilde{\pi}_2\}.$$

Log-linearizing equation (B4) around a non-zero-inflation steady state and assuming symmetry yields

$$\begin{aligned} \left(\frac{\theta}{1 - \beta\mu} + \frac{\omega(1 - \theta)}{1 - \beta\mu(1 + \bar{\pi})} \right) \widehat{p}_t^* &= \left(\frac{\theta}{1 - \beta\mu} + \frac{\omega(1 - \theta)}{1 - \beta\mu(1 + \bar{\pi})} \right) \beta\mu(1 + \bar{\pi}) E_t \widehat{p}_{t+1}^* \\ &+ \left(\frac{\theta}{1 - \beta\mu} + \frac{\omega(1 - \theta)}{1 - \beta\mu(1 + \bar{\pi})} \right) \beta\mu E_t (\pi_{t+1} - \bar{\pi}) \\ &+ \left(\frac{\theta(1 - \beta\mu(1 + \bar{\pi}))}{1 - \beta\mu} + \omega(1 - \theta) \right) \widehat{m}c_t \\ &- \frac{\omega(1 - \theta)\beta\mu}{1 - \beta\mu(1 + \bar{\pi})} (\pi_t - \bar{\pi}), \end{aligned} \quad (\text{B7})$$

which is the analogue of equation (13) in the text.

The Phillips curve associated with this pricing structure can be obtained by combining equations (B3) and (B7) in the usual way. Although the coefficients in the resulting expression are complicated functions of the model parameters, in terms of its general structure, the menu/information costs Phillips curve is given by

$$\pi_t - \bar{\pi} = f(\theta, \omega, \varepsilon, \beta, \bar{\pi}) E_t (\pi_{t+1} - \bar{\pi}) + b(\theta, \omega, \varepsilon, \beta, \bar{\pi}) (\pi_{t-1} - \bar{\pi}) + s(\theta, \omega, \varepsilon, \beta, \bar{\pi}) \widehat{m}c_t,$$

containing both forward- and backward-looking inflation dynamics and having real marginal costs as the driving variable. The relationship between the Phillips curve coefficients, f , b , and s , and the model parameters is complicated but straightforward to analyze numerically.

To this end, with β fixed at 0.99, the four panels in Figure 2 illustrate how f , b , and s vary in response to changes in θ , ω , ε , and $\bar{\pi}$. For this exercise, I set $\theta = 0.50$, $\omega = 0.50$, $\varepsilon = 11$, and $\bar{\pi} = 0.005$ and consider independent variations in each parameter holding the remaining parameters unchanged.

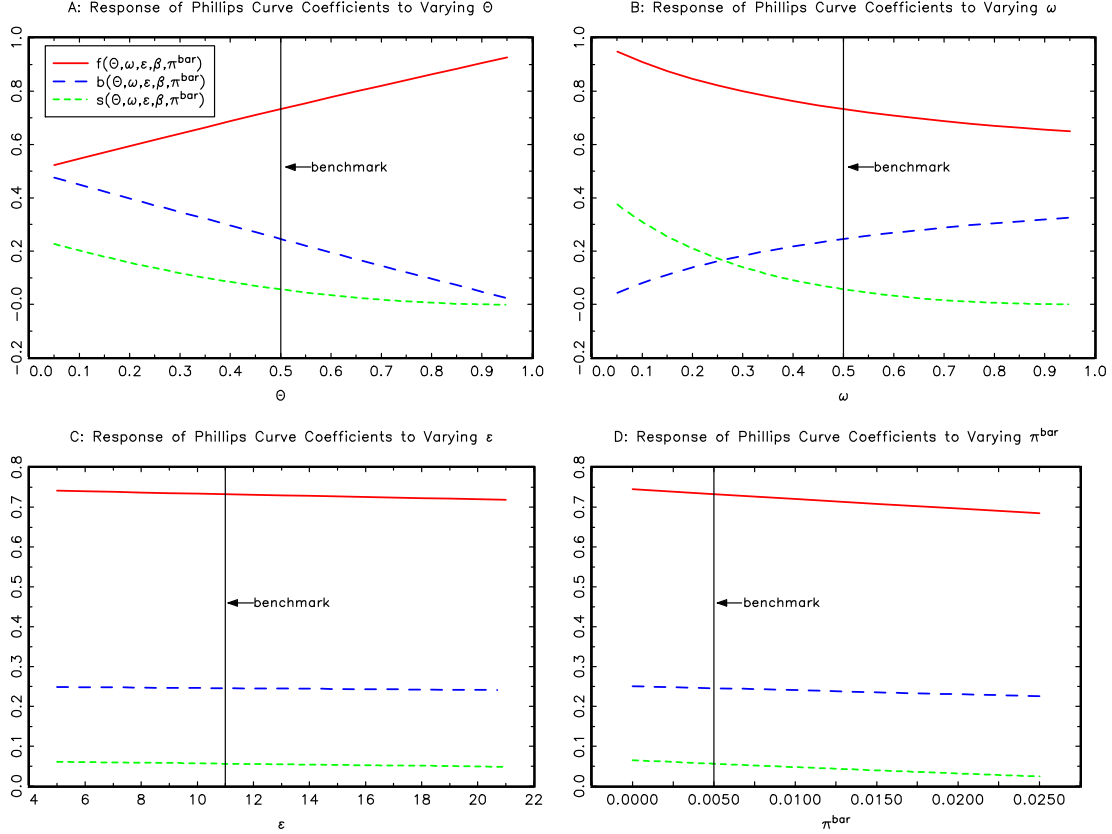


Figure 2: Sensitivity analysis

Figure 2 reveals two important results. First, the model coefficients are sensitive to variation in θ and ω , which implies that the Phillips curve is informative for these two pricing parameters. Second, the model coefficients are not especially sensitive to variation in ε and $\bar{\pi}$, which implies that estimates of θ and ω are likely to be robust to different assumptions regarding ε and $\bar{\pi}$.

C Appendix: The household problem

Households choose consumption, c_t , investment, i_t , their supply of labor, l_t , and their holdings of nominal money balances, m_t , and bonds, b_t , to maximize

$$E_t \sum_{j=0}^{\infty} \beta^j \left[\frac{e^{g_t} (c_{t+j} - \gamma C_{t+j-1})^{1-\sigma}}{1-\sigma} + \frac{\left(\frac{m_{t+j}}{P_{t+j}}\right)^{1-\alpha}}{1-\alpha} - \frac{l_{t+j}^{1+\chi}}{1+\chi} \right], \quad (\text{C1})$$

where $\{\sigma, \alpha, \chi\} \in (0, \infty)$, where $\gamma \in [0, 1)$, and where g_t , $g_t \sim i.i.d. [0, \sigma_g^2]$, is an aggregate consumption-preference shock, subject to the budget constraint

$$c_t + \frac{m_t}{P_t} + \frac{b_t}{P_t} + i_t = w_t l_t + r_t k_t + \frac{(1 + R_{t-1}) b_{t-1}}{P_t} + \frac{m_{t-1}}{P_t} + \frac{\Pi_t}{P_t}$$

and the capital accumulation equation

$$k_{t+1} = (1 - \delta) k_t + i_t,$$

where R_t denotes the nominal interest rate, w_t denotes the consumption real wage, r_t denotes the real rental payment on capital, Π_t denotes the lump-sum profits households earn from dividend payments from firms and the seigniorage revenues households receive from the government, and k_t denotes the capital stock owned by the household. Equation (C1) allows for habit formation, positing that what matters for households is their consumption in relation to lagged aggregate consumption, C_{t-1} .

Since household consumption must always remain above the habit stock $(c_t - \gamma C_{t-1}) > 0$, additive habits are closely related to the notion that there is a subsistence level below which a household's consumption cannot fall. The first-order conditions for the Lagrangian, Λ , associated with the household's problem, include

$$\frac{\partial \Lambda}{\partial c_t} : e^{gt} (c_t - \gamma C_{t-1})^{-\sigma} - \lambda_t = 0, \quad (\text{C2})$$

$$\frac{\partial \Lambda}{\partial l_t} : \lambda_t w_t - l_t^x = 0, \quad (\text{C3})$$

$$\frac{\partial \Lambda}{\partial b_t} : \beta (1 + R_t) E_t \left[\left(\frac{P_t}{P_{t+1}} \right) \lambda_{t+1} \right] - \lambda_t = 0, \quad (\text{C4})$$

$$\frac{\partial \Lambda}{\partial k_{t+1}} : \beta E_t [(r_{t+1} + 1 - \delta) \lambda_{t+1}] - \lambda_t = 0. \quad (\text{C5})$$

Equation (C2) simply defines λ_t , the shadow price of capital, to equal the marginal utility of consumption. Equation (C3) implies that households supply labor up to the point where the marginal rate of substitution between consumption and leisure equals the consumption real wage, w_t . Equation (C4) shows that the bond market clears at an aggregate stock of zero when the expected change in the shadow price of capital equals the ex ante real interest rate. Finally, equations (C5) and (C4) imply that in equilibrium households are indifferent between owning bonds and capital.

Combining equations (C2) and (C4), the log-linear consumption Euler equation is

$$\hat{c}_t = \frac{\gamma}{1 + \gamma} \hat{c}_{t-1} + \frac{1}{1 + \gamma} E_t \hat{c}_{t+1} - \frac{(1 - \gamma)}{\sigma} (R_t - E_t \pi_{t+1} - \rho - gt),$$

which is equivalent to equation (20) in the text.

D Appendix: Aggregate real marginal costs

Cost minimization implies that firms rent capital and labor such that

$$\frac{W_t}{P_t} = mc_t(i) \frac{\kappa y_t(i)}{l_t(i)},$$

implying that a firms' real marginal costs depend on the ratio of the consumption real wage to its marginal productivity of labor, i.e.,

$$mc_t(i) = \frac{1}{\kappa} \frac{w_t l_t(i)}{y_t(i)}.$$

Of course, since all firms face the same rental prices for capital and labor and are subject to the same aggregate technology shock, they employ capital and labor in the same ratio and share the same real marginal costs. Therefore,

$$mc_t(i) = \frac{1}{\kappa} \frac{w_t l_t}{y_t}. \quad (\text{D1})$$

Log-linearizing equation (D1) implies

$$\widehat{mc}_t \simeq \widehat{w}_t + \widehat{l}_t - \widehat{y}_t. \quad (\text{D2})$$

Equation (D2) establishes that, to a first-order log-linear approximation, aggregate real marginal costs depend on the consumption real wage and the aggregate marginal productivity of labor.

The firm-level production function is given by,

$$y_t \equiv \left[\int_0^1 y_t(i)^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}} = \left[\int_0^1 \left([e^{u_t} l_t(i)]^\kappa k_t(i)^{1-\kappa} \right)^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}},$$

which, when log-linearized yields

$$\widehat{y}_t \simeq u_t + \widehat{l}_t + (1 - \kappa) \left(\widehat{k}_t - \widehat{u}_t - \widehat{l}_t \right). \quad (\text{D3})$$

To proceed further, I consider the limiting case in which $\kappa \uparrow 1$ and the role of capital in production tends to zero. Under this assumption the log-linearized resource constraint is $\widehat{c}_t = \widehat{y}_t$. Combining equations (D2) and (D3) with (a log-linearized) equation (C3), the expression for real marginal costs becomes

$$\widehat{mc}_t = \chi \widehat{y}_t - (1 + \chi) u_t - \widehat{\lambda}_t.$$

Now log-linearizing equation (C2) yields

$$\widehat{\lambda}_t = -\frac{\sigma}{(1 - \gamma)} (\widehat{c}_t + \gamma \widehat{c}_{t-1}) + g_t,$$

implying that real marginal costs equal

$$\widehat{mc}_t = \left[\chi + \frac{\sigma}{(1 - \gamma)} \right] \widehat{c}_t - \frac{\sigma \gamma}{(1 - \gamma)} \widehat{c}_{t-1} - (1 + \chi) u_t - g_t,$$

which is equation (21) in the text.

References

- [1] Altig, D., Christiano, L., Eichenbaum, M., and J. Linde, (2004), “Firm-Specific Capital, Nominal Rigidities, and the Business Cycle,” Northwestern University, manuscript.
- [2] Amato, J., and T. Laubach, (2003), “Estimation and Control of an Optimization-Based Model with Sticky Prices and Wages,” *Journal of Economic Dynamics and Control*, 27, pp1181-1215.
- [3] Bilal, M., and P. Klenow, (2004), “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112, pp947-985.
- [4] Blinder, A., (1994), “On Sticky Prices: Academic Theories Meet the Real World,” in Mankiw, N. G. (ed), *Monetary Policy*, University of Chicago Press, Chicago, pp117-150.
- [5] Calvo, G., (1983), “Staggered Contracts in a Utility-Maximising Framework,” *Journal of Monetary Economics*, 12, pp383-398.
- [6] Carlton, D., (1986), “The Rigidity of Prices,” *American Economic Review*, 76, 4, pp637-658.

- [7] Cecchetti, S., (1986), "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines," *Journal of Econometrics*, 31, pp255-274.
- [8] Chang, Y., and S-B. Kim, (2005), "On the Aggregate Labor Supply," *Federal Reserve Bank of Richmond Economic Quarterly*, 91, pp21-37.
- [9] Cho, S., and A. Moreno, (2005), "A Small-Sample Study of the New Keynesian Macro Model," *Journal of Money, Credit, and Banking*, 38, 6, pp1461-1481.
- [10] Christiano, L., Eichenbaum, M., and C. Evans, (2005), "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, pp1-45.
- [11] Clarida, R., Galí, J., and M. Gertler, (1998), "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, February, pp147-180.
- [12] Clarida, R., Galí, J., and M. Gertler, (1999), "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 4, pp1661-1707.
- [13] Clarida, R., Galí, J., and M. Gertler, (2000), "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, February, pp147-180.
- [14] Dixit, A., and J. Stiglitz, (1977), "Monopolistic Competition and Optimum Product Diversity," *American Economic Review Papers and Proceedings*, 67, pp297-308.
- [15] Dennis, R., (2006a), "The Policy Preferences of the U.S. Federal Reserve," *Journal of Applied Econometrics*, 21, pp55-77.
- [16] Dennis, R., (2006b), "Consumption-Habits in a New Keynesian Business Cycle Model," Federal Reserve Bank of San Francisco, manuscript.
- [17] Duffy, J., and P. McNelis, (2001), "Approximating and Simulating the Stochastic Growth Model: Parameterized Expectations, Neural Networks, and the Genetic Algorithm," *Journal of Economic Dynamics and Control*, 25, pp1273-1303.
- [18] Eichenbaum, M., and J. Fisher, (2004), "Evaluating the Calvo Model of Sticky Prices," National Bureau of Economic Research Working Paper #10617.
- [19] Eichenbaum, M., and J. Fisher, (2007), "Estimating the Frequency of Price Reoptimization in Calvo-Style Models," *Journal of Monetary Economics*, forthcoming.
- [20] Erceg, C., Henderson, D., and A. Levin, (2000), "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46, pp281-313.
- [21] Estrella, A., and J. Fuhrer, (2002), "Dynamic Inconsistencies: Counterfactual Implications of a Class of Rational-Expectations Models," *American Economic Review*, 92, 4, pp1013-1028.
- [22] Favero, C., and R. Rovelli, (2003), "Macroeconomic Stability and the Preferences of the Fed. A Formal Analysis, 1961-98," *Journal of Money, Credit, and Banking*, 35, pp545-556.
- [23] Fuhrer, J., (2000), "Optimal Monetary Policy in a Model with Habit Formation," *American Economic Review*, 90, 3, pp367-390.
- [24] Galí, J., and M. Gertler, (1999), "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics*, 44, pp195-222.

- [25] Galí, J., Gertler, M., and D. López-Salido, (2001), “European Inflation Dynamics,” *European Economic Review*, 45, pp1237-1270.
- [26] Gelfand, A., and D. Dey, (1994), “Bayesian Model Choice: Asymptotics and Exact Calculations,” *Journal of the Royal Statistical Society Series B*, 56, pp501-514.
- [27] Gelman, A., (1995), “Inference and Monitoring Convergence,” in Gilks, W., Richardson, S., and D. Spiegelhalter, (eds), *Practical Markov Chain Monte Carlo*, Chapman and Hall Press, London, pp131-143.
- [28] Geweke, G., (1992), “Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments,” in Bernardo, J., Dawid, A., and A. Smith (eds), *Bayesian Statistics*, Clarendon Press, Oxford, 4, pp641-649.
- [29] Geweke, G., (1999), “Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication,” *Econometric Reviews*, 18, pp1-126.
- [30] Giannoni, M., and M. Woodford, (2003), “How Forward-Looking is Optimal Monetary Policy?” *Journal of Money, Credit, and Banking*, 35, 6, pp1425-1469.
- [31] Jung, Y., and T. Yun, (2005), “Implications of Inventory Models for the New Keynesian Phillips Curve,” Kyunghee University, manuscript.
- [32] Kashyap, A., (1995), “Sticky Prices: New Evidence from Retail Catalogs,” *Quarterly Journal of Economics*, February, pp245-274.
- [33] Kim, J., (2000), “Constructing and Estimating a Realistic Optimizing Model of Monetary Policy,” *Journal of Monetary Economics*, 45, pp329-359.
- [34] Laubach, T., and J. Williams, (2003), “Measuring the Natural Rate of Interest,” *Review of Economics and Statistics*, 85, 4, pp1063-1070.
- [35] Levin, A., Onatski, A., Williams, J., and N. Williams, (2006), “Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models,” in Gertler, M., and K. Rogoff (eds), *NBER Macroeconomics Annual 2005*, MIT Press, Cambridge, Massachusetts, pp229-287.
- [36] Liu, Z., and L. Phaneuf, (2005), “Technology Shocks and Labor Market Dynamics: Some Evidence and Theory,” *Journal of Monetary Economics*, forthcoming.
- [37] Lucas, R., (1976), “Econometric Policy Evaluation: A Critique,” *Carnegie-Rochester Conference Series on Public Policy*, 1, pp19-46.
- [38] Mankiw, N. G., (1985), “Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly,” *Quarterly Journal of Economics*, May, pp529-537.
- [39] Mankiw, N. G., and R. Reis, (2002), “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117, 4, pp1295–1328.
- [40] McCallum, B., and E. Nelson, (1999), “Nominal Income Targeting in an Open-Economy Optimizing Model,” *Journal of Monetary Economics*, 43, pp553-578.
- [41] Nakamura, E., and J. Steinsson, (2006), “Five Facts About Prices: A Reevaluation of Menu Cost Models,” Harvard University, mimeo.
- [42] Rabanal, P., and J. Rubio-Ramírez, (2005), “Comparing New Keynesian Models of the Business Cycle: A Bayesian Approach,” *Journal of Monetary Economics*, 52, pp1151-1166.

- [43] Ravenna, F., and C. Walsh, (2006), "Optimal Monetary Policy and the Cost Channel," *Journal of Monetary Economics*, 53, 2, pp199-216.
- [44] Roberts, J., (1995), "New Keynesian Economics and the Phillips Curve" *Journal of Money, Credit, and Banking*, 27, 4, pp975-984.
- [45] Rotemberg, J., (1982), "Sticky Prices in the United States," *Journal of Political Economy*, 60, pp1187-1211.
- [46] Rotemberg, J., and M. Woodford, (1997), "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," in Bernanke, B., and J. Rotemberg, (eds) *NBER Macroeconomics Annual 1997*, MIT Press, Cambridge.
- [47] Rudd, J., and C. Whelan, (2006), "Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?" *American Economic Review*, 96, 1, pp303-320.
- [48] Sahuc, J-G., (2004), "Partial Indexation and Inflation Dynamics: What do the Data Say?" *Applied Economics Letters*, 11, pp827-832.
- [49] Sbordone, A., (2002), "Prices and Unit Labor Costs: A New Test of Price Stickiness," *Journal of Monetary Economics*, 49, pp265-292.
- [50] Sims, C., (2003), "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, pp665-690.
- [51] Smets, F., (2003), "Maintaining Price Stability: How Long is the Medium Term," *Journal of Monetary Economics*, 50, pp1293-1309.
- [52] Smets, F., and R. Wouters, (2005), "Comparing Shocks and Frictions in US and Euro Area Business Cycles: A Bayesian DSGE Approach," *Journal of Applied Econometrics*, 20, pp161-183.
- [53] Taylor, J., (1980), "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88, pp1-24.
- [54] Walsh, C., (2003), "Speed Limit Policies: The Output Gap and Optimal Monetary Policies," *American Economic Review*, 93, 1, pp265-278.
- [55] Woodford, M., (2003), *Interest and Prices*, Princeton University Press, Princeton, New Jersey.
- [56] Zbaraki, M., Ritson, M., Levy, D., Dutta, S., and M. Bergin, (2004), "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economics and Statistics*, 86, 2, pp514-533.
- [57] Zbaraki, M., Levy, D., and M. Bergin, (2007), "The Anatomy of a Price Cut: Discovering Organizational Sources of the Costs of Price Adjustment," Manuscript.