



Reserve Bank of Australia

RESEARCH
DISCUSSION
PAPER

**Dynamic Pricing and
Imperfect Common
Knowledge**

Kristoffer Nimark

RDP 2007-12

DYNAMIC PRICING AND IMPERFECT COMMON KNOWLEDGE

Kristoffer Nimark

Research Discussion Paper
2007-12

December 2007

Economic Research Department
Reserve Bank of Australia

The author thanks Klaus Adam, Leon Berkelmans, Giuseppe Bertola, Giancarlo Corsetti, Fabio Ghironi, James Hansen, Jarkko Jääskelä, Mariano Kulish, Tim Robinson, Thomas Sargent, Michael Woodford, Stephen Wright and seminar participants at the ECB, the European University Institute, Birkbeck College, University of Cambridge, and Nuffield College for valuable comments on earlier drafts. The views expressed in this paper are those of the author and not necessarily those of the Reserve Bank of Australia.

Author: [nimarkk at domain rba.gov.au](mailto:nimarkk@domain.rba.gov.au)

Economic Publications: ecpubs@rba.gov.au

Abstract

This paper introduces private information into the dynamic pricing decision of firms in an otherwise standard new Keynesian model by adding an idiosyncratic component to firms' marginal costs. The model can then replicate two stylised facts about price changes: aggregate inflation responds gradually and with inertia to shocks, while at the same time price changes of individual goods can be quite large. The inertial behaviour of inflation is driven by privately informed firms strategically 'herding' on the public information contained in the observations of lagged aggregate variables. The model also matches the average duration between price changes found in the data and it nests the standard new Keynesian Phillips curve as a special case. To solve the model, the paper derives an algorithm for solving a class of dynamic models with higher-order expectations.

JEL Classification Numbers: E30, F41

Keywords: higher-order expectations, idiosyncratic marginal cost, price dynamics, new Keynesian Phillips curve

Table of Contents

1.	Introduction	1
2.	Idiosyncratic Marginal Costs	4
2.1	The Optimal Reset Price with Imperfect Common Knowledge	5
2.2	Two Limit Cases without Private Information	7
2.2.1	Common marginal costs	8
2.2.2	Large variance of idiosyncratic marginal cost component	8
3.	A Simple General Equilibrium Model	10
3.1	The Model	11
3.2	Expectations and Common Knowledge of Rationality	13
3.3	Solving the Model	15
3.3.1	Output	16
3.3.2	Inflation	17
3.4	Finding a Fixed Point	18
4.	Price Dynamics	19
4.1	Inflation Dynamics and the Size of the Idiosyncratic Shocks	20
4.2	The Model and Actual Inflation Dynamics	22
4.3	The Model and the Evidence on Individual Price Changes	24
5.	Conclusions	27
	Appendix A: The Phillips Curve with Imperfect Common Knowledge	30
	Appendix B: The Law of Motion of the Expectations Hierarchy	32
	References	34

DYNAMIC PRICING AND IMPERFECT COMMON KNOWLEDGE

Kristoffer Nimark

1. Introduction

In standard new Keynesian models, firms set prices to equal a mark-up over expected marginal cost. The real marginal cost is determined by both exogenous and endogenous factors, where the exogenous factors are assumed to be common among all firms. In this paper I relax the assumption of only common exogenous factors by introducing an idiosyncratic component to firms' marginal costs. This does not only improve the realism of the model, but can also help reconcile two apparently conflicting stylised facts that the standard model cannot account for: aggregate inflation responds gradually and with inertia to shocks, while at the same time price changes of individual goods can be quite large.

The inability of the baseline new Keynesian model to match the inertia of inflation is well documented and has spurred economists to suggest explanations, often involving some type of mechanical indexation to past prices.¹ For instance, Galí and Gertler (1999) suggest that a fraction of firms set the price of their own good equal to the previous period's average reset price plus the lagged inflation rate, while Christiano, Eichenbaum and Evans (2005) let a fraction of firms increase the price of their own good according to the lagged inflation rate. Both of these explanations of inflation inertia are attractive since they admit relatively parsimonious representations of realistic inflation dynamics, but they can be criticised as being *ad hoc*. In this paper the inertial behaviour of inflation is driven by *optimising* price-setters.

Private information is introduced into the price-setting problem of the firm through the idiosyncratic component of marginal costs. The optimal price of an individual good depends positively on a firm's own marginal cost and the price chosen by other firms, but individual firms cannot observe the marginal cost of other firms and, therefore, do not know the current price chosen by other firms with certainty.

¹ See, for instance, Fuhrer and Moore (1995), Galí and Gertler (1999) and Galí, Gertler and López-Salido (2001, 2005).

This set-up may be referred to as firms having *imperfect common knowledge* in an environment with strategic complementarities.² In such an environment, it is a well established result that agents tend to put too much weight on public relative to private information.³ In the present model this takes the form of firms ‘herding’ on the publicly observable lagged aggregate variables, inflation and output. This creates the appearance of inflation being partly backward-looking in spite of the fact that all firms are rational and forward-looking.

The idiosyncratic component in firms’ marginal costs also helps to explain that individual price changes are significantly larger than average aggregate price changes. Obviously, increasing the variance of the idiosyncratic component of marginal costs will increase the variance of individual price changes, but this direct effect is not the only one. A firm’s own marginal cost provides a signal about the marginal costs faced by other firms and a large idiosyncratic variance makes this signal less precise. The less precise signal mutes the response of prices to aggregate shocks, since more of a given shock will be attributed to idiosyncratic sources. Increasing the variance of the idiosyncratic component then unambiguously increases the relative magnitude of individual price changes as compared to aggregate price changes.

The idea that incomplete adjustment of prices to aggregate shocks can be explained by information imperfections is not new, but dates back to the Phelps-Lucas island model of the 1970s.⁴ Recently, this idea has had something of a revival. Mankiw and Reis (2002) and Woodford (2002) show how limited information availability, or limited information processing capacities, can produce persistent real effects of nominal disturbances.⁵ Sims (2003) and Mackowiak and Wiederholt (2007) use information processing capacity constraints to explain the inertial responses of aggregate time series to shocks. The model presented here differs from these studies in some important respects that are worth emphasising.

² For example, see Woodford (2002) and Adam (2006).

³ See Morris and Shin (2002) and Chamley (2004).

⁴ See Phelps (1970) and Lucas (1972, 1973, 1975).

⁵ Variants of the Woodford (2002) framework include Hellwig (2005), Adam (2006) and Amato and Shin (2006).

First, through the Calvo mechanism of price adjustment, the model presented here can be made consistent with observed average price durations.⁶ The importance of this assumption boils down to whether one believes that the price stickiness that can be observed in the data causes firms to be forward-looking in a quantitatively important way. In the present model, expectations of future inflation will play a prominent role in determining today's inflation since there is a positive probability that a firm's price may be effective for more than one period. In the papers by Mankiw and Reis (2002), Woodford (2002), and Mackowiak and Wiederholt (2007), the price-setting problem of the firm is a series of static decisions since there is no need for the firm to forecast the future when prices are changed in every period. The dynamic structure of the pricing problem in the present paper makes existing solution methods for models with private information and strategic interaction non-applicable, and we derive a new algorithm to solve the model.⁷ This may be of independent interest.

Second, the models of Mankiw and Reis, Woodford, and Mackowiak and Wiederholt are all closed by using a constant-velocity-of-money type of equation. Here, a richer (but still small) general equilibrium model where households choose how much to consume and how much labour to supply is presented. The model is also more explicit in terms of what firms observe. While the model is too simple to be used to quantify the degree of information imperfections, being explicit prevents us from treating the precision of firms' information as a completely free parameter.

Section 2 derives a Phillips curve under the assumptions of imperfect common knowledge and Calvo pricing. Section 2 also discusses two limit cases of marginal cost structures that preclude any private information, to illustrate how idiosyncratic components in firms' marginal costs can introduce delayed responses to aggregate shocks. Section 3 presents the general equilibrium model and defines the concept of *hierarchies of expectations* and the assumptions that will be imposed on these to solve the model. Section 3 also shows how the recursive structure of the Phillips curve and the IS equation can be exploited to find the solution of the model. Section 4 contains the main results of the paper and

⁶ See Aucremanne and Dhyne (2004), Bils and Klenow (2004), Alvarez *et al* (2005), Klenow and Kryvtsov (2005) and Nakamura and Steinsson (2007).

⁷ See Woodford (2002) or Morris and Shin (2002) for solutions of static decision models.

demonstrates that the model can explain the observed inertia of inflation as well as the observed, relatively large, changes of individual goods prices as compared to the average aggregate price changes, while matching the average duration of prices found in the data. Section 5 concludes.

2. Idiosyncratic Marginal Costs

In most (and perhaps all interesting) economies, one agent's optimal decision depends on the decisions of others. In an economy where all firms and agents are symmetric and all exogenous disturbances are common across firms and agents, knowing the actions of others is a trivial task. An agent can, by observing his own exogenous disturbance, infer the disturbances faced by everybody else and take action based on that information knowing that, in equilibrium, all agents will choose the same action. This is not possible in an economy with idiosyncratic exogenous shocks. Instead, each agent has to form an expectation of the other agents' actions based on what he can observe directly and on *collected* information. The expectation will be imperfect if the collection process adds noise to the observation or if it takes time.

In this paper, these ideas are applied to the price-setting problem of firms that are subject to idiosyncratic marginal cost shocks and where the aggregate price level is only observable with a lag. Individual firms care about the aggregate price level since demand for their own good depends on its price relative to other goods, but due to the idiosyncratic marginal cost shocks, firms cannot infer the aggregate price level perfectly by observing their own marginal cost. The lagged observation then becomes important as a source of information that individual firms use to form expectations about the aggregate price level. The positive correlation between the optimal current price and the lagged price level causes inflation to appear to react to shocks with inertia.

Idiosyncratic marginal cost shocks introduce private information into the price-setting problem of firms and it is demonstrated below how this forces firms to form higher-order expectations, that is, expectations of others' expectations, about marginal cost and future inflation. The variance of the idiosyncratic component of marginal cost determines how accurate a firm's own marginal cost is as an indicator of the average economy-wide marginal cost. By studying the model under two extreme assumptions about this variance, this section also

demonstrates analytically how idiosyncratic marginal cost shocks can introduce delayed responses of inflation to aggregate shocks.

2.1 The Optimal Reset Price with Imperfect Common Knowledge

Apart from the introduction of the idiosyncratic marginal cost component, the framework below is a standard new Keynesian set-up with sticky prices and monopolistic competition. As in Calvo (1983), there is a constant probability $(1 - \theta)$ that a firm, indexed by $j \in (0, 1)$, resets its price in any given period and firms operate in a monopolistically competitive environment. In what follows, all variables are in log deviations from steady-state values. The price level follows

$$p_t = \theta p_{t-1} + (1 - \theta) p_t^* \quad (1)$$

where p_t^* is the average price chosen by firms resetting their price in period t

$$p_t^* = \int p_t^*(j) dj. \quad (2)$$

Firm j 's optimal reset price is the familiar discounted sum of firm j 's expected future nominal marginal costs

$$p_t^*(j) = (1 - \beta\theta) E_t(j) \left[\sum_{i=0}^{\infty} (\beta\theta)^i (p_{t+i} + mc_{t+i}(j)) \right] \quad (3)$$

where β is the firm's discount factor and $E_t(j) [\cdot] \equiv E[\cdot | I_t(j)]$ is an expectations operator conditional on firm j 's information set at time t , which, in this section, is given by:

$$I_t(j) = \left\{ mc_s(j), p_{s-1}, \beta, \theta, \sigma_\varepsilon^2, \sigma_v^2 \mid s \leq t \right\} \quad (4)$$

(for a derivation of the optimal reset price, Equation (3), see Woodford 2003 and the references therein). The structural parameters $\{\beta, \theta, \sigma_\varepsilon^2, \sigma_v^2\}$ and the lagged price level p_{s-1} are common knowledge. The actual economy-wide marginal cost cannot be directly observed (not even with a lag), but firm j can observe its own marginal cost $mc_t(j)$ which is a sum of the economy-wide component mc_t and the firm-specific component $\varepsilon_t(j)$

$$\begin{aligned} mc_t(j) &= mc_t + \varepsilon_t(j) \\ \varepsilon_t(j) &\sim N(0, \sigma_\varepsilon^2) \quad \forall j \in (0, 1). \end{aligned} \quad (5)$$

Since the common and the idiosyncratic components are not distinguishable by direct observation, firm j cannot know with certainty the economy-wide average marginal cost mc_t . The average marginal cost matters for the optimal price of firm j though, since average marginal cost partly determines the current price level. If the average marginal cost process is persistent, then current average marginal cost will also be informative about future marginal costs, and future price levels. To set the price of its good optimally, firm j thus has to form an expectation of the average marginal cost.

The filtering problem faced by the individual firm is thus similar to that faced by the inhabitants of the market ‘islands’ in the well-known Lucas (1975) paper, but with some differences. In Lucas’s model, information is shared among agents between periods so that all agents have the same prior expectation about the aggregate price change, while in this model no such information sharing occurs. This means that since all firms solve a similar signal extraction problem before they set prices, it also becomes relevant for each firm to form higher-order expectations, that is, expectations of average expectations, and so on. By repeatedly substituting in the expression for the price level (1) and the expression for the average reset price (2) into the optimal reset price (3), current inflation can be written as a function of average higher-order expectations of the current average marginal cost and future inflation:

$$\begin{aligned} \pi_t &= (1 - \theta)(1 - \beta\theta) \sum_{k=0}^{\infty} (1 - \theta)^k mc_{t|t}^{(k)} \\ &+ \beta\theta \sum_{k=0}^{\infty} (1 - \theta)^k \pi_{t+1|t}^{(k+1)}. \end{aligned} \quad (6)$$

(The Phillips curve (6) is derived in Appendix A.) The following notation is used to denote higher-order expectations:

$$\begin{aligned} x_t^{(0)} &\equiv x_t \\ x_{t|s}^{(1)} &\equiv \int E[x_t | I_s(j)] dj \\ x_{t|s}^{(2)} &\equiv \int E[x_{t|s}^{(1)} | I_s(j)] dj \\ x_{t|s}^{(k)} &\equiv \int E[x_{t|s}^{(k-1)} | I_s(j)] dj. \end{aligned}$$

In the Phillips curve (6), estimates of order k are weighted by $(1 - \theta)^k$. Since $(1 - \theta)$ is smaller than unity, the impact of expectations is decreasing as the order of expectation increases. This fact is exploited later in order to find a finite dimensional representation of the state of the model. Also, note that $(1 - \theta)$ is decreasing in θ , that is, higher-order expectations are less important when prices are very sticky; when fewer firms change their prices in a given period (when θ is large), average expectations are less important for the firms that actually do change prices.

2.2 Two Limit Cases without Private Information

By the argument presented above, an individual firm needs to form an expectation of the economy-wide average marginal cost (and higher-order estimates of the average marginal cost) to set the price of its own good optimally. To do so, the individual firm uses its knowledge of the structure of the economy, the observations of the lagged price level and of its own marginal cost. The size of the variance of the idiosyncratic component relative to the size of the variance of the average marginal cost innovation determines how accurate firms' estimates will be. Two limiting cases of this variance ratio can help provide some intuition. When the variance of the idiosyncratic component is set to zero, Equation (6) reduces to the standard new Keynesian Phillips curve. In the second, and opposite case, the variance of the idiosyncratic component is assumed to be very large, and this will demonstrate how imperfect information introduces a link between past and current inflation. Both cases preclude any private information, and hence admit analytical solutions. In this section, the simplifying assumption is also made that average marginal cost is driven by the exogenous AR(1) process

$$\begin{aligned} mc_t &= \rho mc_{t-1} + v_t \\ v_t &\sim N(0, \sigma_v^2). \end{aligned} \tag{7}$$

This will facilitate the exposition, and the next section presents a simple general equilibrium model where marginal costs are determined by both exogenous and endogenous factors.

2.2.1 Common marginal costs

If the variance of the idiosyncratic component of firms' marginal costs is equal to zero, that is, $\sigma_\varepsilon^2 = 0$, it follows that:

$$mc_t(j) = mc_t : \forall j. \quad (8)$$

Since firms know the structure of the economy, Equation (8) implies that there is no uncertainty of any order. Formally,

$$mc_t^{(k)} = mc_t : k = 0, 1, 2, \dots, \infty. \quad (9)$$

Since all orders of current marginal cost expectations coincide, so do all orders of future inflation expectations and Equation (6) is reduced to the standard new Keynesian Phillips curve:

$$\pi_t = \beta E_t \pi_{t+1} + \frac{(1-\theta)(1-\beta\theta)}{\theta} mc_t \quad (10)$$

where inflation is completely forward-looking, with marginal cost as the driving variable. By repeated forward substitution, Equation (10) can be written as

$$\pi_t = \frac{(1-\theta)(1-\beta\theta)}{\theta} (1-\rho\beta)^{-1} mc_t \quad (11)$$

which shows that inflation is only as persistent as marginal cost when the individual firm's own marginal cost is a perfect indicator of the economy-wide average.

2.2.2 Large variance of idiosyncratic marginal cost component

This section illustrates the consequences for inflation dynamics when the observation of a firm's own marginal cost holds no information about the economy-wide average. This is strictly true only when the variance of the idiosyncratic marginal cost component reaches infinity, but shocks with infinite variance prevent the law of large numbers from being evoked to calculate the average marginal cost. For illustrative purposes I will temporarily give up on some mathematical rigour. In the following example, the variance of the idiosyncratic component of a firm's marginal cost is 'large enough' for the firm to discard its own marginal cost as an indicator of the economy-wide average. Instead, each firm

uses only the common observation of the lagged price level to form an imperfect expectation of the economy-wide average marginal cost. In this setting, it can be shown that the observation of the lagged price level p_{t-1} perfectly reveals lagged average marginal cost mc_{t-1} . As there is no other source of information available about current average marginal cost, the first-order expectation $mc_t^{(1)}$ is simply given by ρmc_{t-1} . This structure is common knowledge and implies that there is some first-order uncertainty about average marginal costs, that is, $mc_t^{(1)} \neq mc_t$, but no higher-order uncertainty, so that $mc_t^{(k)} = mc_t^{(l)} = \rho mc_{t-1}$ for $k, l > 0$.

Current inflation can be written as a function of actual and first-order expectation of current marginal cost by exploiting the fact that such an expression must nest the solved full information Phillips curve (11) if, by chance, actual and first-order expectation of marginal cost coincide so that $mc_t = mc_t^{(1)}$. From the Phillips curve (6) we know that the coefficient on the actual marginal cost is $(1 - \theta)(1 - \beta\theta)$. To find the coefficient on the first-order expectation of marginal cost, we simply subtract $(1 - \theta)(1 - \beta\theta)$ from the coefficient in the full information solution (11) to get

$$\begin{aligned} \pi_t &= (1 - \theta)(1 - \beta\theta)mc_t \\ &+ \frac{(1 - \theta)(1 - \beta\theta)}{\theta} \left[(1 - \rho\beta)^{-1} - \theta \right] mc_{t|t}^{(1)}. \end{aligned} \quad (12)$$

Using the fact that $mc_t = \rho mc_{t-1} + v_t$ and that $mc_{t|t}^{(1)} = \rho mc_{t-1}$, we can re-arrange Equation (12) into a moving-average representation in the innovations v_t

$$\begin{aligned} \pi_t &= (1 - \theta)(1 - \beta\theta)v_t \\ &+ \frac{(1 - \theta)(1 - \beta\theta)}{\theta} (1 - \rho\beta)^{-1} \sum_{s=1}^{\infty} \rho^s v_{t-s}. \end{aligned} \quad (13)$$

The impulse response to a shock to the average marginal costs will then be hump-shaped if the coefficient on the current innovation v_t is smaller than the coefficient on the lagged innovation v_{t-1} . This will be the case when the persistence parameter ρ is sufficiently large. Thus, the MA representation (13) tells us that the lagged price level will appear to have a positive impact on current inflation only if the average marginal cost follows a persistent process, since only then will lagged inflation hold any information about the marginal costs currently faced by other firms and of future marginal costs. If there is no persistence in marginal costs

($\rho = 0$), lagged inflation does not hold any information relevant to the price-setting problem of the firm and inflation becomes a white noise process.

The assumption of very large idiosyncratic marginal cost shocks, and that it is common knowledge that all firms condition on the same information, made an analytical expression for inflation possible. In the general case, when $0 < \sigma_\varepsilon^2 < \infty$, neither the lagged price level nor the observation of a firm's own marginal cost completely reveal the average marginal cost or other firms' estimates of average marginal cost. Both the firm's own marginal cost as well as the lagged price level will then be needed to form optimal higher-order expectations of marginal costs and, due to the Calvo mechanism, higher-order expectations of future inflation.

3. A Simple General Equilibrium Model

This section presents a simple general equilibrium model where marginal cost is determined by both endogenous and exogenous factors and describes how the model can be solved. The economy consists of households who supply labour and consume goods, firms that produce differentiated goods and set prices, and a central bank that sets the nominal interest rate. Households are subject to economy-wide shocks to their (dis)utility of supplying labour.⁸ The labour supply shock is not directly observable by firms but influences the marginal cost of production. In addition to the labour supply shock and the level of production, firms' marginal costs are also affected by firm-specific wage bargaining shocks and firms cannot, by direct observation, distinguish between the economy-wide labour supply shock and the idiosyncratic bargaining shock. By the same logic as in the previous section, firms then have to form higher-order expectations of the average marginal cost in order to set prices optimally.

This section also formalises the assumption that rational expectations are common knowledge, which simply means that firms and households do not make systematic mistakes given their information sets, and that all firms and households know that all firms and households know, and so on, that all firms and households form rational expectations. This assumption will impose sufficient structure on higher-order expectations to allow the model to be solved.

⁸ Such a shock is estimated in a full information setting in Smets and Wouters (2003).

3.1 The Model

In what follows, lower-case letters denote log deviations from steady-state values of the corresponding capital letter. The representative household maximises

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\exp(\delta_t) \frac{C_t^{1-\gamma}}{1-\gamma} - \exp(\lambda_t) \frac{N_t^{1+\varphi}}{1+\varphi} \right) \quad (14)$$

where: N_t is the aggregate labour supply in period t ; δ_t is a demand shock with zero mean and variance σ_δ^2 ; and β is the discount rate. C_t is the usual CES consumption aggregator

$$C_t = \left(\int_0^1 C_t(j)^{\frac{\vartheta-1}{\vartheta}} dj \right)^{\frac{\vartheta}{\vartheta-1}} \quad (15)$$

and λ_t is a shock to the disutility of supplying labour which is the sum of a persistent component ξ_t and a transitory component η_t :

$$\lambda_t = \xi_t + \eta_t \quad (16)$$

$$\eta_t \sim N(0, \sigma_\eta^2). \quad (17)$$

The persistent component follows an AR(1) process

$$\xi_t = \rho \xi_{t-1} + v_t \quad (18)$$

$$v_t \sim N(0, \sigma_v^2). \quad (19)$$

Firm j produces the differentiated good $Y_t(j)$ using a linear technology with labour as the sole input

$$Y_t(j) = N_t(j). \quad (20)$$

The absence of a storage technology and imposing market clearing implies that aggregate consumption will equal aggregate production

$$Y_t = C_t \quad (21)$$

where the standard CES aggregator is used again. The Euler equation of the representative household then implies the IS-equation

$$y_t = E_t [y_{t+1}] - \frac{1}{\gamma} (i_t - E_t [\pi_{t+1}]) + \delta_t \quad (22)$$

where δ_t is a demand shock with zero mean and variance σ_δ^2 , and i_t is the nominal interest rate. The normative question of how policy should respond to shocks when firms have private information is interesting and is treated by Adam (2006) and Lorenzoni (2007). The focus here is not on the role of monetary policy, and we let the short interest rate follow the simple Taylor-type rule

$$i_t = \phi_\pi \pi_t + \phi_y y_t. \quad (23)$$

In the original formulation of the Taylor-rule (Taylor 1993), monetary policy is set as a function of inflation and the output *gap*, rather than actual output. The slightly different form of the rule (23) is motivated by modelling convenience; by letting the interest rate respond to the same variables that firms can observe, it is not necessary to include the interest rate in firms filtering problem since it does not hold independent information about the labour supply shock λ_t (which partly determines the output gap). The results below should be robust to different formulations of the monetary policy rule, as long as the interest rate does not reveal the labour supply shock λ_t perfectly.

The marginal cost of firm j is the real wage paid at firm j , which is determined by the labour supply decision of households

$$w_t - p_t - \gamma c_t - \varphi n_t - \lambda_t = 0 \quad (24)$$

and a firm-specific wage bargaining shock $\varepsilon_t(j)$. The bargaining shock introduces an idiosyncratic component to firms' marginal costs. Firm j 's marginal cost is

$$mc_t(j) = (\gamma + \varphi) y_t + \lambda_t + \varepsilon_t(j) \quad (25)$$

since $y_t = c_t = n_t$. Firm j 's marginal cost is thus determined by aggregate output y_t , the labour supply shock λ_t , and the idiosyncratic bargaining shock $\varepsilon_t(j)$. The bargaining shock is meant to capture, in a stylised way, the empirical finding that a significant part of the variation in average wages at the firm level appears to be firm specific and uncorrelated to industry-wide changes (see Martins 2003).

The timing of the model is as follows. First, the labour supply shock λ_t is realised. Then, firms and households bargain over wages, where real wages are contracted in the form

$$w_t(j) - p_t = (\gamma + \varphi) y_t + \omega_t(j) \quad (26)$$

where $\omega_t(j) = \lambda_t + \varepsilon_t(j)$. Firms cannot, by direct observation, distinguish between the economy-wide shock to labour supply and the firm-specific bargaining shock. They only observe the sum of the two, $\omega_t(j)$, as well as the component dependent on output, $(\gamma + \varphi)y_t$. The latter can be interpreted as a contract specifying higher hourly wages for (aggregate) overtime. Firms set prices before production takes place and they do not know their own marginal cost with certainty when prices are chosen, but have to form an expectation of what the aggregate output level will be. They will also need to form higher-order expectations of the current average marginal cost and current and future price levels. When prices are set, households choose labour supply and consumption simultaneously. The interest rate and the demand shock are then realised. It is natural to assume that households know the labour supply shock with certainty, and we further assume that there is no information sharing between households and firms. Firm j 's information set when setting the price in period t is thus defined by

$$I_t(j) = \left\{ \omega_t(j), p_{s-1}, y_{s-1}, \beta, \theta, \gamma, \varphi, \sigma_\varepsilon^2, \sigma_v^2, \sigma_\eta^2, \sigma_\delta^2 \mid s \leq t \right\}. \quad (27)$$

3.2 Expectations and Common Knowledge of Rationality

In the two extreme examples in the previous section, firms had no private information and firms' first- and higher-order expectations of marginal cost thus coincided. This is not true in the general case, so first- and higher-order expectations must be treated as separate objects. The fundamental process driving marginal cost in the model is the unobservable economy-wide labour supply shock λ_t , and firms need to form higher-order expectations about this process to set prices optimally. Due to Calvo-pricing, the price-setting decision is forward-looking and firms therefore need to form separate expectations (and higher-order expectations) of the persistent labour supply shock component ξ_t and the transitory component η_t . To simplify notation, the two components of the labour supply shock are collected in the vector \mathbf{x}_t

$$\mathbf{x}_t \equiv [\xi_t \quad \eta_t]'. \quad (28)$$

Common knowledge of rational expectations imposes sufficient structure on expectations to solve the model. Before formalising this notion, the concept of a *hierarchy of expectations* needs to be defined.

Definition 1: Let firm j 's hierarchy of expectations of \mathbf{x}_t from order l to m be the vector

$$\mathbf{x}_{t|t}^{(l:m)} \equiv \left[\mathbf{x}_{t|t}^{(l)} \quad \mathbf{x}_{t|t}^{(l+1)} \quad \dots \quad \mathbf{x}_{t|t}^{(m-1)} \quad \mathbf{x}_{t|t}^{(m)} \right]'. \quad (29)$$

In the solution strategy, the hierarchy of expectations of current labour supply shock is treated as the ‘fundamental’ variable, or the state, of the model. We want to be able to write any order of expectation of the endogenous variables – inflation and output – as functions of the hierarchy from order zero to infinity. To this end, the following assumption is imposed on higher-order expectations.

Assumption 1: It is common knowledge that agents’ expectations are rational (model consistent). Let $\mathcal{M} : \mathbb{R}^\infty \rightarrow \mathbb{R}^\infty$ be a mapping from the hierarchy of expectations of \mathbf{x}_t in period t to the expected hierarchy of expectations in period $t+1$.

$$E \left[\mathbf{x}_{t+1|t+1}^{(0:\infty)} \mid \mathbf{x}_{t|t}^{(0:\infty)} \right] \equiv \mathcal{M} \left(\mathbf{x}_{t|t}^{(0:\infty)} \right). \quad (30)$$

Common knowledge of rational expectations then implies that

$$E \left[\mathbf{x}_{t+1|t+1}^{(0:\infty)} \mid \mathbf{x}_{t|t}^{(k:\infty)} \right] = \mathcal{M} \left(\mathbf{x}_{t|t}^{(k:\infty)} \right) \quad \forall k \geq 0. \quad (31)$$

Let $\mathcal{T} : \mathbb{R}^\infty \rightarrow \mathbb{R}$ be a mapping from the hierarchy of expectations of \mathbf{x}_t in period t to the endogenous variable z_t

$$z_t = \mathcal{T} \left(\mathbf{x}_{t|t}^{(0:\infty)} \right). \quad (32)$$

Common knowledge of rational expectations then implies that

$$E \left[z_{t+1} \mid \mathbf{x}_{t|t}^{(k:\infty)} \right] = \mathcal{T} \left(\mathcal{M} \left(\mathbf{x}_{t|t}^{(k:\infty)} \right) \right) \quad \forall k \geq 0. \quad (33)$$

Assumption 1 is a natural generalisation of the assumption of rational expectations in a common information setting to the private information case.⁹ The mapping \mathcal{M} represents the actual law of motion for the contemporaneous expectations hierarchy. The first part of Assumption 1 simply states that firms use the actual law of motion of the hierarchy to form expectations of future values of the hierarchy

⁹ The full information rational expectations assumption is nested in Assumption 1. To see this, set $\mathbf{x}_{t|t}^{(k)} = \mathbf{x}_{t|t}^{(l)} \forall k, l \geq 0$ and let \mathcal{M} be the exogenous process (16) and \mathcal{T} the function that maps the state into an endogenous outcome.

and that this is common knowledge. The second part makes the same statement about expectations of variables that are functions of the hierarchy of labour supply shock expectations.

For something to be common knowledge, it is not enough that it is commonly believed, it must also be true. Setting $k = 1$ in Equations (31) and (33) makes firms' expectations rational. That Equations (31) and (33) apply to all $k \geq 0$ makes it *common* knowledge, so that all firms know that all firms know, and so on, that all firms have rational expectations.

Since the model is linear, the mappings \mathcal{M} and \mathcal{T} will be linear functions. This means that it does not matter whether Assumption 1 is imposed directly on average expectation hierarchies or the assumption is on individual firms' expectations before taking averages. The practical purpose of Assumption 1 is the same as the standard rational expectations assumption in full information models: it allows all terms involving inflation expectations to be substituted out from the Phillips curve (6). Inflation can then be expressed as a function of the state of the model, that is, the hierarchy of expectations of the current labour supply shock.

3.3 Solving the Model

The model is solved by an iterative version of the method of undetermined coefficients. Conjecture (and verify in Appendix B) that the hierarchy of labour supply shock expectations follows the vector autoregression:

$$\mathbf{x}_{t|t}^{(0:\infty)} = M\mathbf{x}_{t-1|t-1}^{(0:\infty)} + N\mathbf{v}_t \quad (34)$$

where

$$\mathbf{v}_t = [\mathbf{v}_t \quad \eta_t \quad \delta_t]'. \quad (35)$$

The hierarchy of expectations $\mathbf{x}_{t|t}^{(0:\infty)}$ is the state of the model and in Appendix B it is demonstrated that Assumption 1 provides enough structure on higher-order expectations to find the law of motion (34). The main intuition behind the method is that the actual, or zero-order, expectation is given exogenously. The first-order expectation is pinned down by being a rational expectation of the zero-order expectation. Common knowledge of rationality can then be applied to recursively determine the law of motion for higher-order expectation, so that the second-order expectation is a rational expectation of the first-order expectation, the third-order

expectation is a rational expectation of the second-order expectation, and so on. The Kalman filter plays a dual role in this process. Not only is it used by firms to estimate the average expectation hierarchy, but since this hierarchy is made up of the average of the very same estimates, it will also determine the law of motion of the hierarchy, that is, determine the matrices M and N in the conjectured law of motion (34).¹⁰

For a given M in the law of motion (34), output and inflation can be found as functions of the current state of the expectation hierarchy of the labour supply shock \mathbf{x}_t . We want a solution in the following form

$$\pi_t = \mathbf{c}\mathbf{x}_{t|t}^{(0:\infty)} \quad (36)$$

$$y_t = \mathbf{d}\mathbf{x}_{t|t}^{(0:\infty)} + \delta_t \quad (37)$$

in which case the dynamics of inflation and output are completely characterised by Equations (36) and (37) together with the law of motion (34).

3.3.1 Output

Households know the labour supply shock with certainty and form rational expectations about future output and expected real interest rates. Together with the conjectured form of the solved model, Equations (34)–(37) allow the output Euler Equation (22) to be rewritten as

$$\mathbf{d}\mathbf{x}_t^{(0:\infty)} + \delta_t = \mathbf{d}M\mathbf{x}_{t|t}^{(0:\infty)} - \frac{1}{\gamma} \left(\phi_\pi \mathbf{c}\mathbf{x}_{t|t}^{(0:\infty)} + \phi_y \mathbf{d}\mathbf{x}_{t|t}^{(0:\infty)} - \mathbf{c}M\mathbf{x}_{t|t}^{(0:\infty)} \right) + \delta_t \quad (38)$$

where the Taylor-type rule (23) was used to substitute out the nominal interest rate. The fact that households know the actual labour supply shock with certainty means that expected output and the real interest rate are functions of the complete hierarchy of expectations. Matching coefficients in Equation (38) implies that the vector \mathbf{d} must satisfy the identity

$$\mathbf{d} = \mathbf{d}M - \frac{1}{\gamma} (\phi_\pi \mathbf{c} + \phi_y \mathbf{d} - M\mathbf{c}) \quad (39)$$

¹⁰ The Kalman filter plays a similar dual role in Woodford (2002).

3.3.2 Inflation

In the model, prices are set before output is realised, and since marginal cost depends on aggregate output, firms have to form an expectation of aggregate output. We can use the rationality assumption and the marginal cost function (25) to get firm j 's expectations of its own marginal cost

$$E[mc_t(j) | I_t(j)] \equiv (\gamma + \varphi) E[y_t | I_t(j)] + \omega_t(j). \quad (40)$$

Taking averages across firms yields an expression for the average expectation of firms' own marginal cost

$$\widehat{mc}_t^{(0)} = (\gamma + \varphi) y_t^{(1)} + \lambda_t \quad (41)$$

since $\int \omega_t(j) = \lambda_t$. Invoking common knowledge of rational expectations yields a general expression for a k -order expectation of firms' marginal cost

$$\widehat{mc}_t^{(k)} = (\gamma + \varphi) y_t^{(k+1)} + \lambda_t^{(k)}. \quad (42)$$

Using the conjectured law of motion for the hierarchy of expectations (34) and inflation (36) and Equation (42) to write all terms in the Phillips curve (6) as functions of the expectation hierarchy of \mathbf{x}_t implies

$$\begin{aligned} \mathbf{c}\mathbf{x}_{t|t}^{(0:\infty)} &= (1 - \theta)(1 - \beta\theta) \sum_{k=0}^{\infty} (1 - \theta)^k \left(\mathbf{d}\mathbf{x}_{t|t}^{(k:\infty)} + \mathbf{1}_{1 \times 2} \mathbf{x}_{t|t}^{(k)} \right) \\ &+ \beta\theta \sum_{k=0}^{\infty} (1 - \theta)^k \mathbf{c}M\mathbf{x}_{t|t}^{(k+1:\infty)} \end{aligned} \quad (43)$$

since $\lambda_t = \mathbf{1}_{1 \times 2} \mathbf{x}_t$. Equating coefficients implies that the vector \mathbf{c} must satisfy

$$\mathbf{c} = \mathbf{a}D + \mathbf{b} \begin{bmatrix} 0 & \mathbf{c}M \\ 0 & 0 & \mathbf{c}M \\ \vdots & & \ddots & \mathbf{c}M \\ 0 & \dots & \dots & 0 \end{bmatrix} \quad (44)$$

where

$$D = (\gamma + \varphi) \begin{bmatrix} \mathbf{0}_{1 \times 2} & \mathbf{d} & & \\ \mathbf{0}_{1 \times 2} & \mathbf{0}_{1 \times 2} & \mathbf{d} & \\ \vdots & & \ddots & \vdots \\ \mathbf{0}_{1 \times 2} & \dots & \dots & \mathbf{0}_{1 \times 2} \end{bmatrix} + \begin{bmatrix} \mathbf{1}_{1 \times 2} & \mathbf{0} & & \\ \mathbf{0}_{1 \times 2} & \mathbf{1}_{1 \times 2} & \mathbf{0} & \\ \vdots & & \ddots & \vdots \\ \mathbf{0}_{1 \times 2} & \dots & \dots & \mathbf{1}_{1 \times 2} \end{bmatrix}. \quad (45)$$

The row vectors \mathbf{a} and \mathbf{b} are given directly by Equation (43)

$$\mathbf{a} = \left[(1 - \theta)(1 - \beta\theta) \quad (1 - \theta)^2(1 - \beta\theta) \quad \cdots \quad (1 - \theta)^\infty(1 - \beta\theta) \right] \quad (46)$$

$$\mathbf{b} = \left[\beta\theta \quad \beta\theta(1 - \theta) \quad \cdots \quad \beta\theta(1 - \theta)^\infty \right]. \quad (47)$$

3.4 Finding a Fixed Point

Solving the model implies finding a fixed point for \mathbf{c} , \mathbf{d} , M and N . The derivations above involve expectations of up to infinite order, which is problematic since in practise we cannot solve the model using infinite dimensional vectors and matrices. To obtain an approximation that can be made arbitrarily accurate, we exploit the fact that the impact of expectations is decreasing as the order of expectation increases. Intuitively, the magnitude of a price-setter's response to a unit change in his expectation of marginal cost or future inflation is decreasing as the order of expectation increases. In Phillips curve (6) this can be seen from the fact that the term raised to the power of the order of expectation k , $(1 - \theta)^k$, is smaller than one. As k becomes large, this term approaches zero. Together with the fact that the unconditional variance of expectations cannot increase as the order of expectation increases, an arbitrarily accurate solution can be found by including a sufficiently large but finite number of orders of expectations in the state of the model.¹¹

In practise, the model is solved by guessing a candidate number k^* of how many orders of expectations to include. A fixed point for the model with $\mathbf{x}_{t|t}^{(0:k^*)}$ as the state vector can then be found by direct iteration on Equations (38) and (44) and the expression for M and N in Appendix B. After a solution has been found, we check whether adding one more order of expectations and re-solving the model, with $\mathbf{x}_{t|t}^{(0:k^*+1)}$ as the state, changes the impact of a shock to marginal cost on inflation enough to motivate, including higher orders of expectations. Once we are satisfied with the accuracy of the solution, we can simulate the model using Equations (34), (36) and (37).

¹¹ That the variance of higher-order expectations cannot increase with the order of the expectation is implied by common knowledge of rationality. To see why, define a k^{th} -order expectation error as $e_t^{(k)} \equiv x_{t|t}^{(k)} - x_{t|t}^{(k-1)}$. The error $e_t^{(k)}$ must be orthogonal to $x_{t|t}^{(k)}$ if $x_{t|t}^{(k)}$ is a rational expectation of $x_{t|t}^{(k-1)}$. The fact that $\text{var}(e_t^{(k)}) + \text{var}(x_{t|t}^{(k)}) = \text{var}(x_{t|t}^{(k-1)})$ and that variances are non-negative yields the desired result.

4. Price Dynamics

This section presents the main results of the paper. By simulating the model described in Section 3, a number of issues are addressed. First, it is demonstrated how the variance of the firm-specific wage-bargaining shock influences the degree of inflation inertia, and that the model can explain the positive coefficient found on lagged inflation in estimates of the hybrid new Keynesian Phillips curve. Second, it is demonstrated that the model can replicate the observed large magnitudes of individual price changes as compared to changes in the aggregate price level, as well as match the observed average duration of individual prices found in the data.

Throughout this section, some of the parameters of the model will be held fixed. These are the persistence of the labour supply shock ρ , the discount factor β , the inverse of the intertemporal elasticity of consumption γ , the curvature of the disutility of supplying labour φ and the parameters in the Taylor-type rule. These will be set as $\{\rho, \beta, \gamma, \varphi, \phi_\pi, \phi_y\} = \{0.9, 0.995, 2, 2, 1.5, 0.5\}$. The choice of the exogenous persistence parameter roughly reflects the persistence of various measures of marginal cost (for instance, the labour share in GDP). The precise value of ρ is not important, though some exogenous persistence is necessary to generate interesting results. The parameterisation of the discount factor β at 0.995 reflects that a period in the model should be interpreted as being one month in order to make the assumption that inflation is observed with a one period lag realistic. The Calvo parameter θ should thus be interpreted as the fraction of firms that do not change prices in a given month and, unless otherwise stated, it will be set to $\theta = 0.9$, which implies an average price duration of 10 months. The exact choice of the parameters in the period utility function and the Taylor-type rule (γ , φ , ϕ_π and ϕ_y), are not crucial for the results below.

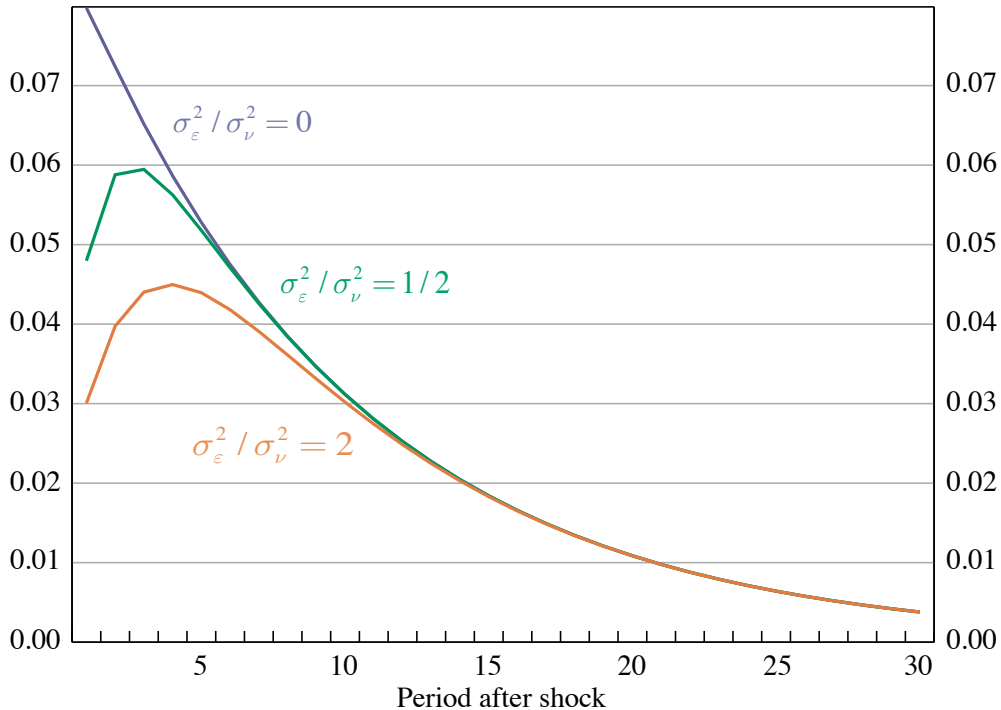
The focus in the analysis below is on the role played by the idiosyncratic shocks to firms' marginal costs. The variance of these shocks relative to the variance of economy-wide disturbances determines how precise a firm's own marginal cost is as an indicator of the economy-wide averages. As a matter of normalisation, the variance of all other shocks are held fixed, as the variance of the idiosyncratic marginal cost shocks is varied in the exercises below. The variance of the demand shock σ_δ^2 and the variance of the innovation to the persistent labour supply shock component σ_v^2 are therefore set to unity.

The labour supply shock λ_t is a compound of two shocks with different persistence. The technical reason for this set-up is that if it was a process with a single innovation in each period, the lagged price level would perfectly reveal the lagged labour supply shock and any information-induced dynamics of inflation would be short-lived. Since the focus of the paper is the consequences of private information, rather than imperfect information in general, the variance of the transitory labour supply shock component σ_η^2 is set to 0.01. This is small enough to make the effects of firms' confounding persistent and transitory economy-wide shocks insignificant, but large enough so that lagged inflation does not completely reveal the persistent labour supply shock component ξ_t .

4.1 Inflation Dynamics and the Size of the Idiosyncratic Shocks

Figure 1 illustrates how the variance σ_ε^2 of the idiosyncratic wage bargaining shocks affects the dynamic response of inflation to a unit shock to the persistent component ξ_t of the labour supply shock λ_t . The blue curve is the response with the idiosyncratic marginal cost shock variance set to zero. The green and orange curves are the impulse responses with the idiosyncratic shock variance set to $1/2$ and 2 , respectively.

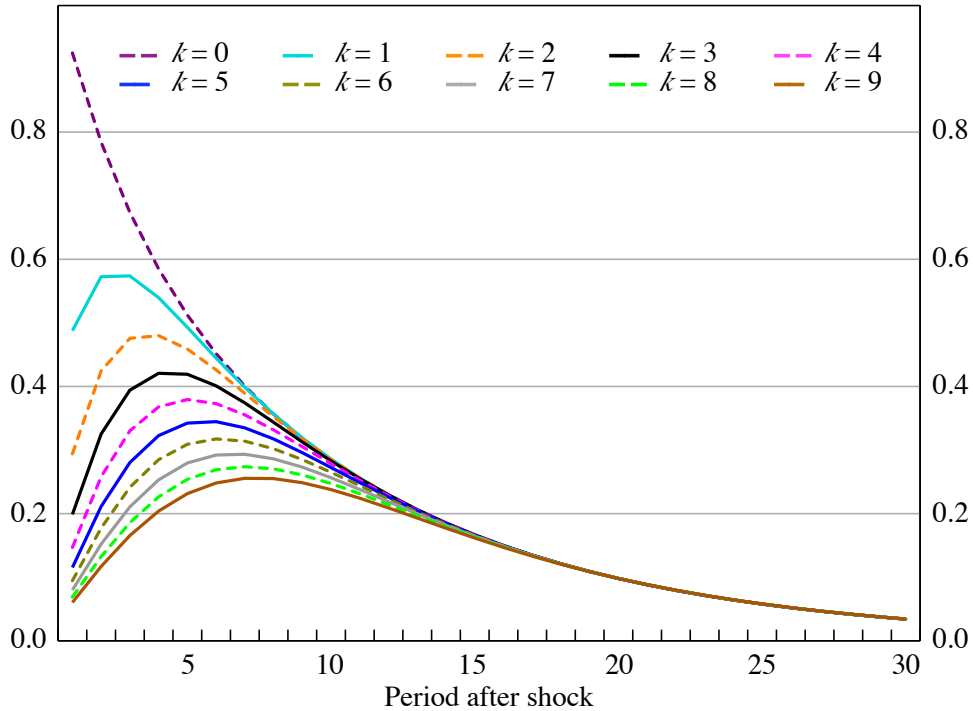
Figure 1: Impulse Response of Inflation to Unit Average Marginal Cost Shock



Three things are worth pointing out. First, with a zero idiosyncratic component variance, the model replicates the full information response, with monotonic convergence to the mean after the shock. Second, with a non-zero variance of the idiosyncratic marginal cost component, the response of inflation is hump-shaped, with the peak of the hump appearing later the larger is the variance σ_ε^2 . Third, the larger this ratio is, the smaller is the first-period impact of a marginal cost shock and the lower is inflation at the peak. Since the underlying labour supply shock in all three cases decreases monotonically, and in a shape identical to the inflation response with a zero variance ratio, the humps must be driven by the dynamics of the higher-order expectations of average marginal cost. Figure 2 displays the dynamics of the hierarchy of marginal cost expectations up to the third order, that is, $\widehat{mc}_t^{(0:3)}$, after a unit shock to the persistent component ξ_t of the labour supply shock with the variance of the idiosyncratic marginal cost shock set to $\sigma_\varepsilon^2 = 2$.

Figure 2 shows that the average first-order expectation moves less than the zero-order marginal cost on impact. The idiosyncratic component thus works as ‘noise’

Figure 2: Impulse Response of Marginal Cost Hierarchy to Unit Average Marginal Cost Shock
 $\sigma_\varepsilon^2 = 2$



in the filtering problem that smooths out estimates of the innovations in the average marginal cost process. In addition, higher-order estimates move less on impact than first-order estimates. The key to understanding the dynamics of the higher-order estimates is that common knowledge of rationality implies a recursive relationship between orders of expectations. Firms' first-order estimate of average marginal cost is rational given their information set, but they know that due to firms confounding idiosyncratic and economy-wide disturbances, shocks are underestimated on average (that is, average first-order expectations move less than the actual shock). Therefore, for a given change in first-order expectations on impact, second-order expectations move less as firms expect other firms to, on average, underestimate the shock in the impact period. A similar argument applies to third- and higher-order expectations and explains why the response of expected marginal cost in the impact period is decreasing in the order of expectation.

In the model presented here, individual firms have two different types of information: the private observation of its own marginal cost and the public observation of lagged inflation and output. A well-established result from the literature on social learning is that when agents receive both private and public signals, and there are strategic complementarities in actions, agents tend to put too much weight on the public signal relative to its precision, which is often referred to as 'herding'.¹² Herding slows down the social learning process by making the endogenous signal, that is, the observed aggregate behaviour, less informative about the underlying exogenous shock. Another way to understand the inertia of inflation is thus that firms herd on public signals that are only observable with a lag.

4.2 The Model and Actual Inflation Dynamics

In the preceding sections I presented qualitative evidence in the form of hump-shaped impulse responses on how imperfect common knowledge can introduce inflation inertia. In this section, the question is asked whether the model can account for the observed inflation inertia in US, euro area and Australian data, with quantitatively realistic amounts of information imperfections. This question

¹² See Morris and Shin (2002) and Chamley (2004).

is pursued by generating data from the simple general equilibrium model and estimating the hybrid new Keynesian Phillips curve

$$\pi_t = \mu_f E_t \pi_{t+1} + \kappa mc_t + \mu_b \pi_{t-1} \quad (48)$$

by generalised method of moments (GMM). If the model is the true data-generating process, then the hybrid new Keynesian Phillips curve is of course misspecified. The experiment of interest is thus to check whether the misspecified econometric model (48) applied to the theoretical model would produce results similar to those obtained when the hybrid new Keynesian Phillips curve is estimated on actual data. Galí *et al* (2005) provide a range of estimates for the US and the euro area, using slightly different choices of instruments and formulations of the orthogonality condition. The estimates of the backward-looking parameter γ_b ranges from 0.035 to 0.27 for the euro area and from 0.32 to 0.36 for the US. A robust feature across methodologies is that the estimated inflation inertia is lower in the euro area than in the US. Estimating the same equation on Australian data from the post-inflation-targeting period (1993:Q1–2006:Q4) yields backward-looking parameter estimates of between 0.36 and 0.5.

Table 1 displays estimates of the hybrid new Keynesian Phillips curve (48) using simulated data from the model for different ratios of the variance of the bargaining shock over the variance of average marginal cost.¹³

Table 1: Parameter Estimates of the Hybrid New Keynesian Phillips Curve from Simulated Data

$\sigma_\varepsilon^2 / \sigma_{mc}^2$	μ_b	μ_f	κ
1/10	0.19	0.81	0.024
1/2	0.34	0.66	0.013
1	0.41	0.59	0.010

The simulated data was transformed to ‘quarterly’ frequencies by taking three period averages. The orthogonality condition

$$E_t [\pi_t - (1 - \mu_b)\pi_{t+1} - \mu_b\pi_{t-1} - \kappa mc_t] = 0 \quad (49)$$

¹³ Three hundred ‘monthly’ observations were transformed into 100 ‘quarterly’ observations. The estimates displayed in the table are averages over 20 independent samples.

was then estimated by GMM using all available lagged variables, that is, marginal costs, output and inflation rates, as instruments.

Table 1 tells us that the variance of the idiosyncratic bargaining shock only needs to be about 1/2 the size of the variance of the average marginal cost for the model to generate the inflation inertia observed in the US. Quantitative information on the magnitude of unexplained firm-level variations in real wages are hard to come by, but there are some studies where this information can be extracted as a by-product. Martins (2003) investigates the competitiveness of the Portuguese garment industry labour market using yearly data. He finds that between 30 and 40 per cent of average wage variations across firms cannot be explained by labour market conditions, changes in the skills of workers, production techniques or (time dependent) firm-level fixed effects. It is not clear that this is representative of other industries and countries. However, the fact that the model requires a variance ratio of 1/2 to match US inflation inertia cannot be considered conspicuously unrealistic.

4.3 The Model and the Evidence on Individual Price Changes

In addition to matching the observed inertia of aggregate inflation, the model can also reproduce some features of the behaviour of individual goods prices. Two widely cited studies on the frequency of individual good's price changes are Bils and Klenow (2004) on US consumer price data, and that carried out by the Inflation Persistence Network of the central banks within the Eurosystem. The latter is summarised in Alvarez *et al* (2005). Both of these studies find that prices of individual goods change infrequently, but less so for the US than for the euro area. Bils and Klenow report a median probability of a good changing price in a given month of around 25 per cent. The distribution of the frequencies of price changes is not symmetric and the average price duration therefore differs from the reciprocal of the median frequency. The average duration of prices in the US is around 7 months, rather than $(1/0.25) = 4$ months. A recent study by Nakamura and Steinsson (2007), using a more detailed data set than Bils and Klenow, finds that the average price duration of consumer prices in the US is between 8 and 11 months for the sample period 1998–2005. Alvarez *et al* find a median duration of consumer prices in the euro area of about 10 months.

The micro evidence also suggests that individual prices are much more volatile than the aggregate price level. Klenow and Kryvtsov (2005) find that, conditional on a price change occurring, the average absolute individual price change is 8.5 per cent for the US, which is the same number reported by Nakamura and Steinsson (2007). This can be compared with the average absolute monthly change in the US CPI of 0.32 per cent for the period 1981–2005. Average absolute individual price changes are thus about 25 times larger than average absolute aggregate price changes.

The price level is non-stationary in the model of this paper which, together with the Calvo mechanism, prevents an analytical derivation of the absolute average size of individual price changes. What can be done instead is to use the law of motion of the system (34) and the inflation Equation (36) to simulate the model and then compute the average absolute change in both the aggregate price level and of a typical individual good's price. The inflation Equation (36) implies that the price level follows

$$p_t = \mathbf{cx}_{t|t}^{(0:\infty)} + p_{t-1} \quad (50)$$

since $\pi_t = p_t - p_{t-1}$. The optimal reset price of good j in period t is then given by

$$p_t^*(j) = (1 - \theta\beta)\varepsilon_t(j) + (1 - \theta)^{-1}\mathbf{cx}_{t|t}^{(1:\infty)}(j) + p_{t-1} \quad (51)$$

since $p_t = \theta p_{t-1} + (1 - \theta) \int p_t^*(j) dj$. It is clear from Equation (51) that increased price stickiness, that is, a higher θ , will increase the relative size of individual price changes to aggregate price changes; when fewer firms change prices in a given period, they have to move by more for a given change in the aggregate price level to occur. In the model here, the idiosyncratic component affects the relative size of individual and aggregate price changes through two additional and distinct channels.

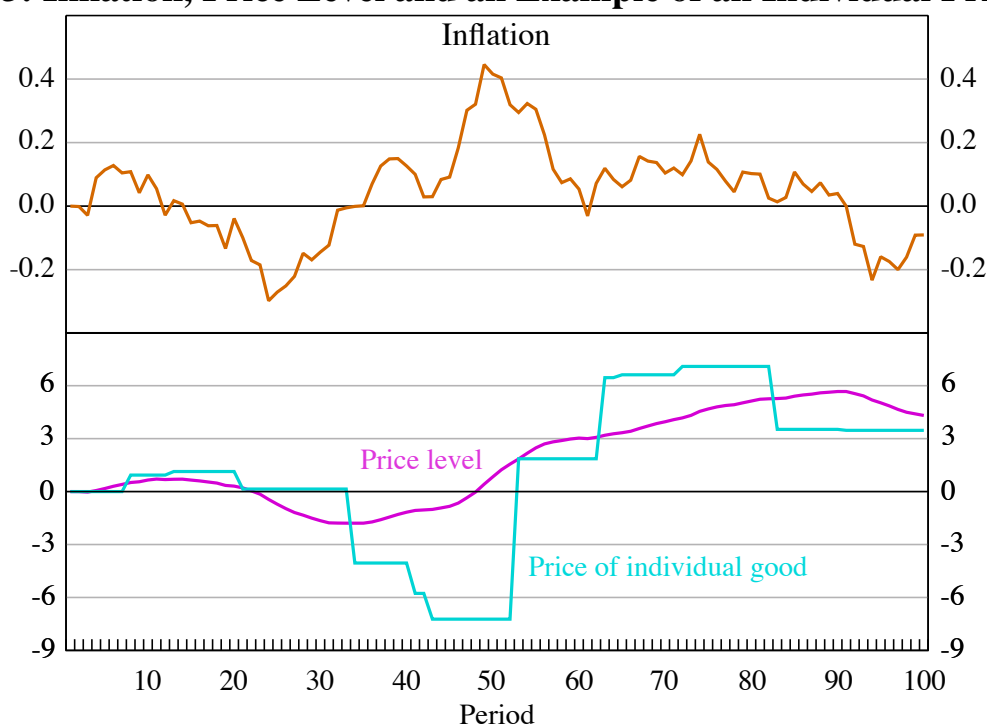
First, there is the direct effect of larger idiosyncratic marginal cost shocks on the magnitude of individual price changes; more volatile individual marginal costs will lead to more volatile individual prices.

Second, the larger is the variance of the idiosyncratic component, the less precise is a firm's own marginal cost as a signal of aggregate marginal cost. As demonstrated above, increasing the variance of the idiosyncratic component leads to more muted responses to aggregate shocks. Since the idiosyncratic shocks cancel out in

aggregate, it is only the second indirect effect that affects the aggregate price level. The magnitude of individual price changes relative to aggregate price changes therefore unambiguously increases as the variance of the idiosyncratic component of marginal costs is increased.

In Figure 3, simulated series of inflation, the aggregate price level and the price of a typical good are plotted. The variance of the idiosyncratic marginal cost shock is set to 1/2 of the variance of average marginal cost. It is clear from the bottom panel of the figure that, conditional on a price change occurring, individual good prices are much more volatile than the aggregate price level.

Figure 3: Inflation, Price Level and an Example of an Individual Price Path



The aggregate price level path and the path of an individual good's price can be simulated for various ratios of idiosyncratic and economy-wide marginal cost shock variances so as to compute the relative size of the average absolute price changes of an individual good compared to the average absolute aggregate price level changes. The results of this exercise for different degrees of price stickiness are reported in Table 2.

Comparing the two columns to the right in Table 2 tells us that the average magnitude of individual price changes relative to aggregate price changes are

Table 2: Relative Magnitude of Individual and Aggregate Price Changes

$\sigma_\varepsilon^2 / \sigma_{mc}^2$	$E \Delta p_t(j) / E \Delta p_t $	
	$\theta = 0.8$	$\theta = 0.9$
1/10	6.3	13.4
1/2	6.6	17.5
1	7.1	20.4

significantly larger when the average price duration is 10 months ($\theta = 0.9$) as compared to an average price duration of 5 months ($\theta = 0.8$), suggesting that it may be important to match the frequency of price changes when calibrating models to replicate the observed average price changes of individual goods. With an average price duration of 10 months, the model can match the numbers reported by Klenow and Kryvtsov (2005) and Nakamura and Steinsson (2007) of individual price changes, on average, being 25 times larger than aggregate price changes, by setting the idiosyncratic marginal cost shock variance to be approximately 1.4 times the variance of the economy-wide average marginal cost.

5. Conclusions

In this paper I have argued that, when firms have idiosyncratic components to their marginal costs, they cannot compute the current price level perfectly before they choose their own optimal price. Instead, firms have to form an estimate of the price level using the information contained in their own marginal cost and in observations of past inflation and output. This structure, coupled with the Calvo mechanism of price adjustment, results in a Phillips curve with a role for higher-order expectations of marginal cost and future inflation. Even though the pricing decision is entirely forward-looking, lagged inflation will have an impact on current inflation since lagged inflation contains information relevant to the optimal price of the firm. This effect is amplified by firms having private information about marginal costs, which induces ‘herding’, or over-weighting, of the information in the publicly-observed lagged aggregate variables. This information effect can explain the positive coefficient found on lagged inflation in estimates of the hybrid new Keynesian Phillips curve.

The idiosyncratic component of marginal costs can also help explain the fact that individual price changes are significantly larger than aggregate price changes. In addition to increasing the volatility of individual marginal costs, a higher variance

of the idiosyncratic component makes it harder for individual firms to filter out the economy-wide component from the observation of their own marginal cost. This second effect reduces the responses of prices to economy-wide shocks and is a similar result to that of Mackowiak and Wiederholt (2007). They find that firms will choose to allocate less attention to aggregate variables when idiosyncratic conditions are very volatile.

The modelling philosophy of Mackowiak and Wiederholt differs from that used in this paper in their more abstract approach. They do not explicitly specify what the idiosyncratic conditions facing the firms in their model are, which makes it harder to judge whether the constraints on information processing capacity that are necessary to match the data are realistic or not. In this paper, the precision of firms' information is determined by the relative variance of individual firms' marginal costs and the economy-wide aggregate marginal cost. This explicit approach makes it possible, at least in principle, to compare variances of the model with variances in the data. It is then also possible to ask whether information imperfections are likely to be large enough in reality to be important for the dynamics of inflation. It is hard to argue that firms' own marginal costs and lagged inflation and output are the only information available to agents in reality. However, a *necessary* condition for explanations of economic phenomena to be plausibly based on limited information availability (or limited capacity to process information) is that quantities that are immediately and costlessly observable to agents are not too informative.

In the specific case considered here, the firm-level idiosyncratic wage variances necessary to replicate US inflation dynamics was about half of the overall variance of real wages. In principle, this could be compared to variance ratios in the data, but in practice, such a comparison is constrained by the limited availability of firm-level data. Martins (2003) is a rare study that provides some information on the relative size of firm-specific and industry-wide variances of wages paid in one industry (garment production) in one country (Portugal). If Martins' numbers are representative, actual firm-specific shocks are more volatile than what is necessary for the model to replicate the observed inflation inertia, but somewhat less volatile than necessary to replicate the observed magnitude of individual price changes. However, given the very limited availability of firm-level data, it is prudent to avoid drawing strong conclusions about what should be considered a realistic lower bound on how imperfectly informed firms can be.

The more explicit approach to information imperfections of this paper also points out directions for further research. The present model suggests an explanation for the observed higher-inflation inertia in the US relative to the euro area. European wage bargaining is often centralised, while in the US a larger fraction of wages are set at the firm level.¹⁴ Hence, there is thus likely to be more firm-level variation in wages in the US than in Europe which, in the model, would lead to more inertia, and could thus explain the observed differences. Comparing the predictions of the model with a larger cross-section of countries may be another way to validate the main mechanism of the model.

Finally, through the Calvo mechanism, the model is consistent with the micro evidence on the average duration of prices, but it also implies that firms need to be forward-looking when they set prices. The fact that agents in the model make dynamic choices, rather than a series of static choices, renders existing solution techniques inapplicable. The model is solved by imposing that rational expectations are common knowledge. This assumption, together with a structural model that implies that the impact of higher-order expectations are decreasing as the order of expectation increases, makes it possible to derive a solution algorithm of arbitrary accuracy. Though some of the details of the algorithm are relegated to Appendix B, it may be of independent interest to some readers.

¹⁴ See OECD (1997).

Appendix A: The Phillips Curve with Imperfect Common Knowledge

The price level follows

$$p_t = \theta p_{t-1} + (1 - \theta) \int p_t^*(j) dj. \quad (\text{A1})$$

The optimal price of firm j is given by

$$p_t^*(j) = (1 - \beta\theta) E_t(j) \sum_{i=0}^{\infty} (\beta\theta)^i (p_{t+i} + mc_{t+i}(j)) \quad (\text{A2})$$

or, equivalently,

$$p_t^*(j) = (1 - \beta\theta) E_t(j) (p_t + mc_t(j)) + E_t(j) \beta\theta p_{t+1}^*(j). \quad (\text{A3})$$

The average reset price p_t^* is then given by

$$p_t^* = (1 - \beta\theta) \left(p_t^{(1)} + mc_t^{(1)} \right) + \beta\theta p_{t+1|t}^* \quad (\text{A4})$$

Repeated substitution of Equations (A1) and (A4) into Equation (A3) and averaging across firms and substituting into (A1) yields

$$\begin{aligned} p_t = & (1 - \theta)(1 - \beta\theta) \sum_{k=0}^{\infty} ((1 - \beta\theta)(1 - \theta))^k mc_t^{(k)} + \\ & + \frac{1}{1 + \beta - \theta\beta} p_{t-1} + (1 - \theta)\theta\beta \sum_{k=0}^{\infty} ((1 - \beta\theta)(1 - \theta))^k p_{t+1|t}^{*(k+1)}. \end{aligned} \quad (\text{A5})$$

Use $p_t = p_{t-1} + \pi_t$ and $p_{t+1}^* = \frac{\pi_t}{(1-\theta)} + p_t$ to get

$$\begin{aligned} \pi_t = & (1 - \theta)(1 - \beta\theta) \sum_{k=0}^{\infty} ((1 - \beta\theta)(1 - \theta))^k mc_t^{(k)} \\ & + \frac{\beta(1 - \theta)}{1 + \beta - \theta\beta} p_{t-1} + \theta\beta \sum_{k=0}^{\infty} ((1 - \beta\theta)(1 - \theta))^k \pi_{t+1|t}^{(k+1)} \\ & + (1 - \theta)\theta\beta \sum_{k=0}^{\infty} ((1 - \beta\theta)(1 - \theta))^k p_{t|t}^{(k+1)} \end{aligned} \quad (\text{A6})$$

which can be rewritten as

$$\begin{aligned}
\pi_t &= (1 - \beta\theta)(1 - \theta) \sum_{k=0}^{\infty} ((1 - \beta\theta)(1 - \theta))^k mc_t^{(k)} \\
&\quad + \beta\theta \sum_{k=0}^{\infty} ((1 - \beta\theta)(1 - \theta))^k \pi_{t+1|t}^{(k+1)} \\
&\quad + \theta\beta(1 - \theta) \left[\sum_{k=0}^{\infty} ((1 - \beta\theta)(1 - \theta))^k \pi_t^{(k+1)} \right].
\end{aligned} \tag{A7}$$

Common knowledge of rationality implies that any order l of average expectations of current inflation can be written as

$$\begin{aligned}
\pi_t^{(l)} &= (1 - \beta\theta)(1 - \theta) \sum_{k=l}^{\infty} ((1 - \beta\theta)(1 - \theta))^k mc_t^{(k)} \\
&\quad + \beta\theta \sum_{k=l}^{\infty} ((1 - \beta\theta)(1 - \theta))^k \pi_{t+1|t}^{(k+1)} \\
&\quad + \theta\beta(1 - \theta) \left[\sum_{k=l}^{\infty} ((1 - \beta\theta)(1 - \theta))^k \pi_t^{(k+1)} \right]
\end{aligned} \tag{A8}$$

which allows us to substitute out all orders of current inflation expectations in Equation (A8) to get Equation (6) in the main text:

$$\begin{aligned}
\pi_t &= (1 - \beta\theta)(1 - \theta) \sum_{k=l}^{\infty} (1 - \theta)^k mc_t^{(k)} \\
&\quad + \beta\theta \sum_{k=l}^{\infty} (1 - \theta)^k \pi_{t+1|t}^{(k+1)}.
\end{aligned} \tag{A9}$$

Appendix B: The Law of Motion of the Expectations Hierarchy

Define a new state vector X_t as

$$X_t \equiv \begin{bmatrix} \mathbf{x}_t^{(0:\infty)} \\ \mathbf{x}_{t-1|t-1}^{(0:\infty)} \end{bmatrix}, \quad X_t = WX_{t-1} + \begin{bmatrix} N \\ 0 \end{bmatrix} \mathbf{v}_t \quad (\text{B1})$$

$$W \equiv \begin{bmatrix} M & 0 \\ I & 0 \end{bmatrix}, \quad \mathbf{v}_t = [v_t \quad \eta_t \quad \delta_t]'. \quad (\text{B1})$$

For a given W , firm j 's estimate of X_t will evolve according to the updating equation

$$X_{t|t}^{(1)}(j) = WX_{t-1|t-1}^{(1)}(j) + K \left[S_t(j) - LX_{t-1|t-1}^{(1)}(j) \right] \quad (\text{B2})$$

where K is the Kalman gain matrix and $S_t(j)$ is firm j 's observation vector. L is a matrix that maps an expected state into an expected observation vector. They are given by

$$S_t(j) = [\omega_t(j) \quad \pi_{t-1} \quad y_{t-1}]' \quad (\text{B3})$$

$$L = \begin{bmatrix} \mathbf{1}_{1 \times 2} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{c} \\ \mathbf{0} & \mathbf{0} & \mathbf{d} \end{bmatrix} \quad (\text{B4})$$

$$K = PL'(LPL' + \Sigma_{\varepsilon\varepsilon})^{-1} \quad (\text{B5})$$

$$P = W(P - PL'(LPL' + \Sigma_{\varepsilon\varepsilon})^{-1}LP)W' + \Sigma_{uu} \quad (\text{B6})$$

$$\Sigma_{\varepsilon\varepsilon} = \begin{bmatrix} \sigma_\varepsilon^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \sigma_y^2 \end{bmatrix}, \quad \Sigma_{uu} = N \begin{bmatrix} \sigma_v^2 & 0 & 0 \\ 0 & \sigma_\eta^2 & 0 \\ 0 & 0 & \sigma_y^2 \end{bmatrix} N'.$$

Take averages of Equation (B2) across firms and use the definitions of the observables in S_t to write the average updating equation as

$$X_{t|t}^{(1)} = (I - KL)WX_{t-1|t-1}^{(1)} + KLWX_{t-1} + KL \begin{bmatrix} N \\ 0 \end{bmatrix} \mathbf{v}_t + KLe_t. \quad (\text{B7})$$

Finally, use the definition of the hierarchy $\mathbf{x}_{t|t}^{(0:\infty)}$:

$$\mathbf{x}_{t|t}^{(0:\infty)} = \begin{bmatrix} \mathbf{x}_t & \mathbf{x}_{t|t}^{(1:\infty)} \end{bmatrix}' \quad (\text{B8})$$

and the fact that the exogenous state $\mathbf{x}_t|_t$ follows the known process

$$\mathbf{x}_t = R\mathbf{x}_{t-1} + [v_t \ \delta_t]', \quad R = \begin{bmatrix} \rho & 0 \\ 0 & 0 \end{bmatrix} \quad (\text{B9})$$

to identify M and N :

$$M = \begin{bmatrix} R & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{2 \times 2} & \mathbf{0} \\ \mathbf{0} & [W - K LW]_{11} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ [K L W]_{11} \end{bmatrix} \quad (\text{B10})$$

$$N = \begin{bmatrix} \mathbf{1}_{1 \times 2} & 0 \\ \mathbf{0} & \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ [K L N]_{11} \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & K_3 \end{bmatrix} \quad (\text{B11})$$

where $[\cdot]_{11}$ is the upper left submatrix of the appropriately partitioned matrix $[\cdot]$.

References

Adam K (2006), ‘Optimal Monetary Policy with Imperfect Common Knowledge’, *Journal of Monetary Economics*, 54(2), pp 267–301.

Alvarez LJ, E Dhyne, MM Hoeberichts, C Kwapil, H Le Bihan, P Lønneman, F Martins, R Sabbatini, H Stahl, P Vermeulen and J Vilmunen (2005), ‘Sticky Prices in the Euro Area: A Summary of New Micro Evidence’, European Central Bank Working Paper Series No 563.

Amato J and H Shin (2006), ‘Imperfect Common Knowledge and the Information Value of Prices’, *Economic Theory*, 27(1), pp 213–241.

Aucremanne L and E Dhyne (2004), ‘How Frequently Do Prices Change? Evidence Based on the Micro Data Underlying the Belgian CPI’, European Central Bank Working Paper Series No 331.

Bils M and PJ Klenow (2004), ‘Some Evidence on the Importance of Sticky Prices’, *Journal of Political Economy*, 112(5), pp 947–985.

Calvo GA (1983), ‘Staggered Prices in a Utility-Maximizing Framework’, *Journal of Monetary Economics*, 12(3), pp 383–398.

Chamley CP (2004), *Rational Herds: Economic Models of Social Learning*, Cambridge University Press, Cambridge.

Christiano LJ, M Eichenbaum and CL Evans (2005), ‘Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy’, *Journal of Political Economy*, 113(1), pp 1–45.

Fuhrer J and G Moore (1995), ‘Inflation Persistence’, *The Quarterly Journal of Economics*, 110(1), pp 127–159.

Galí J and M Gertler (1999), ‘Inflation Dynamics: A Structural Econometric Analysis’, *Journal of Monetary Economics*, 44(2), pp 195–222.

Galí J, M Gertler and JD López-Salido (2001), ‘European Inflation Dynamics’, *European Economic Review*, 45(7), pp 1237–1270.

Galí J, M Gertler and JD López-Salido (2005), ‘Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve’, *Journal of Monetary Economics*, 52(6), pp 1107–1118.

Hellwig C (2005), ‘Heterogeneous Information and the Benefits of Public Information Disclosures’, University of California at Los Angeles Economics Online Papers No 283.

Klenow PJ and O Kryvtsov (2005), ‘State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?’, NBER Working Paper No 11043.

Lorenzoni G (2007), ‘News Shocks and Optimal Monetary Policy’, NBER Working Paper No 12898.

Lucas RE, Jr (1972), ‘Expectations and the Neutrality of Money’, *Journal of Economic Theory*, 4(2), pp 103–124.

Lucas RE, Jr (1973), ‘Some International Evidence on Output-Inflation Tradeoffs’, *American Economic Review*, 63(3), pp 326–334.

Lucas RE, Jr (1975), ‘An Equilibrium Model of the Business Cycle’, *Journal of Political Economy*, 83(6), pp 1113–1144.

Mackowiak BA and M Wiederholt (2007), ‘Optimal Sticky Prices under Rational Inattention’, Centre for Economic Policy Research Discussion Paper No 6243.

Mankiw NG and R Reis (2002), ‘Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve’, *Quarterly Journal of Economics*, 117(4), pp 1295–1328.

Martins PS (2003), ‘Firm Wage Differentials in a Competitive Industry: Some Matched Panel Evidence’, *International Journal of Manpower*, 24(4), pp 336–346.

Morris S and HS Shin (2002), ‘Social Value of Public Information’, *American Economic Review*, 92(5), pp 1521–1534.

Nakamura E and J Steinsson (2007), ‘Five Facts about Prices: A Reevaluation of Menu Cost Models’, Harvard University, mimeo.

OECD (1997), ‘Chapter 3: Economic Performance and the Structure of Collective Bargaining’, in *OECD Employment Outlook*, OECD, Paris, pp 63–92.

Phelps ES (1970), ‘Introduction: The New Microeconomics in Employment and Inflation Theory’, in ES Phelps *et al*, *Microeconomic Foundations of Employment and Inflation Theory*, WW Norton and Company Inc, New York, pp 1–23.

Sims CA (2003), ‘Implications of Rational Inattention’, *Journal of Monetary Economics*, 50(3), pp 665–690.

Smets F and R Wouters (2003), ‘An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area’, *Journal of the European Economic Association*, 1(5), pp 1123–1175.

Taylor JB (1993), ‘Discretion Versus Policy Rules in Practice’, *Carnegie-Rochester Conference Series on Public Policy*, 39, pp 195–214.

Woodford M (2002), ‘Imperfect Common Knowledge and the Effects of Monetary Policy’, in E Phelps, A Alchian and C Holt (eds), *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honour of Edmund S. Phelps*, Princeton University Press, Princeton, pp 25–58.

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