Available online at www.sciencedirect.com
ScienceDirect
Journal of Monetary Economics 55 (2008) 365–382
Journal of MONETARY ECONOMICS
www.elsevier.com/locate/jme

Dynamic pricing and imperfect common knowledge ☆

Kristoffer Nimark*

Research Department, Reserve Bank of Australia, 65 Martin Place, NSW 2001, Sydney, Australia

Received 1 August 2006; received in revised form 30 November 2007; accepted 12 December 2007
Available online 28 December 2007

Abstract

Introducing private information into the dynamic pricing decision of firms by adding an idiosyncratic component to marginal cost can help explain two stylised facts about price changes: Aggregate inflation responds gradually and with inertia to shocks at the same time as individual price changes can be 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: D8; E3

Keywords: Private information; Dynamic price setting; Idiosyncratic marginal cost; Sticky prices; Inflation inertia

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. While convenient from a modelling perspective, this assumption is clearly unrealistic. This paper relaxes the assumption of only common exogenous factors, by introducing an idiosyncratic component in 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 at the same time as price changes of individual goods are 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

☆The author thanks Klaus Adam, Leon Berkelmans, Giuseppe Bertola, Giancarlo Corsetti, Fabio Ghironi, James Hansen, Jarkko Jaaskela, Mariano Kulish, Tim Robinson, Thomas Sargent, Michael Woodford, Stephen Wright and seminar participants at the ECB, the EUI, Birkbeck College, University of Cambridge, and Nuffield College for valuable comments on earlier drafts.
*Tel.: +61 2 9551 8872.
E-mail address: nimarkk@rba.gov.au

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2007.12.008

prices.[1] For instance, Gali 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 et al. (2005) let a fraction of firms increase their own good prices with 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 the present 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. 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 (2005) 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.

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 (2005) 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 (2002), Woodford (2002) and Mackowiak and Wiederholt (2005) are all closed by using a constant-velocity-of-money type of equation. Here, a richer (but still small), general equilibrium model is presented where households choose how much to consume and how much labour to supply. The present model is also more explicit in terms of what firms observe. While the model is too simple

[1] See for instance Fuhrer and Moore (1995), Gali and Gertler (1999) and Galí et al. (2003).

[2] E.g. Woodford (2002) and Adam (2007).

[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 Amato and Shin (2006), Adam (2007) and Hellwig (2004).

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

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

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.

In the next section, a Phillips curve is derived under the assumptions of imperfect common knowledge and Calvo pricing. The next section also uses two limit cases of marginal cost structures that preclude any private information, to illustrate how idiosyncratic components in firms' marginal cost 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 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 (perhaps in 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 we apply these ideas to the price setting problem of a firm that is 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.

The idiosyncratic marginal cost shock introduces private information into the price setting problem of firms and below we show how this forces firms to form higher order expectations, i.e. expectations of other's 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 limit assumption 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 will reset 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:

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

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

$$p_t^* = \int p_t^*(j)\,\mathrm{d}j. \tag{2}$$

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

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

where β is the firm's discount factor and $\mathrm{E}_t(j)[\cdot] \equiv \mathrm{E}[\cdot|I_t(j)]$ an expectations operator conditional on firm j's information set at time t

$$I_t(j) = \{mc_s(j), p_{s-1}, \beta, \theta, \sigma_\varepsilon^2, \sigma_v^2 | s \leqslant t\} \tag{4}$$

(for a derivation of the optimal reset price (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. Actual economy-wide marginal cost cannot be directly observed (not even with a lag), but firm j can observe his own marginal cost $mc_t(j)$ which is a sum of the economy-wide component mc_t and the firm specific component $\varepsilon_t(j)$

$$mc_t(j) = mc_t + \varepsilon_t(j),$$

$$\varepsilon_t(j) \sim \mathrm{N}(0, \sigma_\varepsilon^2) \quad \forall j \in (0,1). \tag{5}$$

Since the common and the idiosyncratic component are not distinguishable by direct observation, firm j cannot know with certainty what the economy-wide average marginal cost mc_t is. 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 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' model, information is shared among agents between periods so that all agents have the same prior about the expected aggregate price change, while in our 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, i.e. 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 (3) current inflation can be written as a function of average higher order expectations of current marginal cost and future inflation[8]

$$\pi_t = (1-\theta)(1-\beta\theta)\sum_{k=0}^{\infty}(1-\theta)^k mc_{t|t}^{(k)} + \beta\theta\sum_{k=0}^{\infty}(1-\theta)^k \pi_{t+1|t}^{(k+1)}, \tag{6}$$

where the following notation for higher order expectations was used

$$x_{t|t}^{(0)} \equiv x_t,$$

$$x_{t|s}^{(1)} \equiv \int \mathrm{E}[x_t|I_s(j)]\,\mathrm{d}j,$$

$$x_{t|s}^{(2)} \equiv \int \mathrm{E}[x_{t|s}^{(1)}|I_s(j)]\,\mathrm{d}j,$$

$$x_{t|s}^{(k)} \equiv \int \mathrm{E}[x_{t|s}^{(k-1)}|I_s(j)]\,\mathrm{d}j.$$

In (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, i.e. higher order expectations are less important when prices are very sticky: When fewer firms change their prices in

[8]The Phillips curve (6) is derived in the Appendix of the working paper version of this article, available at http://www.rba.gov.au/PublicationsAndResearch/RDP/.

a given period, i.e. 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, individual firms need to form an expectation of the economy-wide average marginal cost (and higher order estimates of 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 and 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 limit cases of this variance ratio can help intuition. When the variance of the idiosyncratic component is set to zero, (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 made that average marginal cost is driven by the exogenous AR(1) process

$$mc_t = \rho mc_{t-1} + v_t,$$

$$v_t \sim \mathrm{N}(0, \sigma_v^2). \tag{7}$$

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

2.2.1. Common marginal costs

If we set the variance of the idiosyncratic component of firms' marginal costs equal to zero, i.e. $\sigma_\varepsilon^2 = 0$, it follows that

$$mc_t(j) = mc_t \ : \ \forall j. \tag{8}$$

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

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

Since all orders of current marginal cost expectations coincide, so does all orders of future inflation expectations and (6) is reduced to the standard New-Keynesian Phillips Curve

$$\pi_t = \beta \mathrm{E}_t \pi_{t+1} + \frac{(1-\theta)(1-\beta\theta)}{\theta} mc_t, \tag{10}$$

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

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

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

In this section we illustrate 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 prevents us from invoking the law of large numbers to calculate average marginal cost. For illustrative purposes we will temporarily give up on some mathematical rigor. 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 $\rho mc_{t-1.}$ This structure is common knowledge and implies that there is some first order uncertainty about average marginal costs, i.e. $mc_t^{(1)} \neq mc_t$ but no higher order uncertainty so that $mc_t^{(k)} = mc_t^{(l)} = \rho mc_{t-1} : k, l > 0$.

We can write current inflation as a function of actual and the first order expectation of current marginal cost by exploiting that such an expression must nest the solved full information Phillips curve (11) if, by chance, actual and the 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

$$\pi_t = (1-\theta)(1-\beta\theta)mc_t + \frac{(1-\theta)(1-\beta\theta)}{\theta}[(1-\rho\beta)^{-1} - \theta]mc_{t|t}^{(1)}. \tag{12}$$

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

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

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. The MA representation (13) thus tells us that the lagged price level will appear to have a positive impact on current inflation only if average marginal cost follows a persistent process, since it is only then that lagged inflation holds any information about the marginal costs currently faced by other firms and of future marginal costs. If there is no persistence in marginal costs, that is if ρ is zero, 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 it possible to find an analytical expression for inflation. 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 reveals the average marginal cost or other firms' estimates of average marginal cost. Both the firm's own marginal cost and 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 are formed.

3. A simple general equilibrium model

In this section a simple general equilibrium model is set up where marginal cost is determined by both endogenous and exogenous factors and the method used to solve the model is described. The economy consists of households who supply labour and consume goods, firms that produce differentiated goods and set prices and a monetary policy authority that sets the nominal interest rate. Households are subject to economy-wide shocks to their utility of consumption and (dis)utility of supplying labour.[9] 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 shock 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 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

[9] Such a shock is estimated in a full information setting in Smets and Wouters (2002).

form rational expectations. This assumption will impose sufficient structure on higher order expectations to allow us to solve the model.

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

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

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

$$C_t = \left(\int_0^1 C_t(j)^{(\varepsilon-1)/\varepsilon} \, dj \right)^{\varepsilon/(\varepsilon-1)} \tag{15}$$

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

$$\lambda_t = \xi_t + \eta_t, \tag{16}$$

$$\eta_t \sim N(0, \sigma_\eta^2). \tag{17}$$

The transitory component η_t of the labour supply shock λ_t prevents the lagged price level to perfectly reveal the lagged value of the persistent component ξ_t. The persistent component follows an AR(1) process:

$$\xi_t = \rho \xi_{t-1} + v_t, \tag{18}$$

$$v_t \sim N(0, \sigma_v^2). \tag{19}$$

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

$$Y_t(j) = N_t(j). \tag{20}$$

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

$$Y_t = C_t, \tag{21}$$

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

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

where 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 (2007) 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:

$$i_t = \phi_\pi \pi_t + \phi_y y_t. \tag{23}$$

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 though, 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 intratemporal labour supply decision of households

$$w_t - p_t - \gamma c_t - \varphi n_t - \lambda_t = 0 \tag{24}$$

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

$$mc_t(j) = (\gamma + \varphi)y_t + \lambda_t + \varepsilon_t(j), \tag{25}$$

where we used that $y_t = c_t = n_t$. 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 seem to be firm specific and uncorrelated to industry wide changes (see Martins, 2003).

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

$$w_t(j) - p_t = (\gamma + \varphi)y_t + \omega_t(j), \tag{26}$$

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, but only observe the sum of the two, $\omega_t(j)$, and 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 firms 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 current marginal cost and current and future price levels. When prices are set, households choose labour supply and consumption simultaneously with the determination of the interest rate and the demand shock is 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

$$I_t(j) = \{\omega_t(j), p_{s-1}, y_{s-1}, \beta, \theta, \gamma, \varphi, \sigma_\varepsilon^2, \sigma_v^2, \sigma_\eta^2, \sigma_\varepsilon^2 | s \leqslant t\}. \tag{27}$$

3.2. Expectations and common knowledge of rationality

In the two limit 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, and first and higher order expectations then has to 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 $\mathbf{x}_t$

$$\mathbf{x}_t \equiv [\xi_t \;\; \eta_t]'. \tag{28}$$

We assume common knowledge of rational expectations which imposes sufficient structure on expectations to solve the model. We formalise this notion here, but first we define the concept of a *hierarchy of expectations*.

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

$$\mathbf{x}_{t|t}^{(l:m)} \equiv [\mathbf{x}_{t|t}^{(l)} \;\; \mathbf{x}_{t|t}^{(l+1)} \;\; \cdots \;\; \mathbf{x}_{t|t}^{(m-1)} \;\; \mathbf{x}_{t|t}^{(m)}]'. \tag{29}$$

In our 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. Towards this end, we impose the following assumption on higher order expectations.

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

$$\mathrm{E}[\mathbf{x}_{t+1|t+1}^{(0:\infty)}|\mathbf{x}_{t|t}^{(0:\infty)}]\equiv\mathscr{M}(\mathbf{x}_{t|t}^{(0:\infty)}). \tag{30}$$

Common knowledge of rational expectations then implies that

$$\mathrm{E}[\mathbf{x}_{t+1|t+1}^{(0:\infty)}|\mathbf{x}_{t|t}^{(k:\infty)}]=\mathscr{M}(\mathbf{x}_{t|t}^{(k:\infty)})\quad\forall k\geqslant 0. \tag{31}$$

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

$$z_t=\mathscr{T}(\mathbf{x}_{t|t}^{(0:\infty)}). \tag{32}$$

Common knowledge of rational expectations then implies that

$$\mathrm{E}[z_{t+1}|\mathbf{x}_{t|t}^{(k:\infty)}]=\mathscr{T}(\mathscr{M}(\mathbf{x}_{t|t}^{(k:\infty)}))\quad\forall k\geqslant 0. \tag{33}$$

Assumption 1 is a natural generalisation of the assumption of rational expectations in a common information setting to the private information case.[10] The mapping $\mathscr{M}$ 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 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, that it must also be true. Setting $k=1$ in (31) and (33) makes firms' expectations rational. That (31) and (33) apply to all $k\geqslant 0$ 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 $\mathscr{M}$ and $\mathscr{T}$ will be linear functions. This means that it does not matter whether Assumption 1 is imposed directly on average expectation hierarchies or 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 us to substitute out all terms involving inflation expectations in the Phillips curve (6). Inflation can then be written as a function solely of the state of the model, i.e. 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. We conjecture that the hierarchy of labour supply shock expectations follows the vector autoregression[11]

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

where

$$\mathbf{v}_t=[v_t\ \ \eta_t\ \ \varepsilon_t]'. \tag{35}$$

The hierarchy of expectations $\mathbf{x}_{t|t}^{(0:\infty)}$ is the state of the model and in the Appendix we show how 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.

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

[11]That the law of motion of the hierarchy takes exactly this form is verified in the Appendix of the working paper version of this article, available at http://www.rba.gov.au/PublicationsAndResearch/RDP/ where more details on how to solve the model also can be found.

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 (34).[12]

For a given M in the law of motion (34) we can find output and inflation 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:

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

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

so that the dynamics of inflation and output are completely characterised by (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 (34)–(37) this allows us to rewrite the output Euler equation (22) as

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

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

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

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

$$\mathrm{E}[mc_t(j)|I_t(j)] \equiv (\gamma + \varphi)\mathrm{E}[y_t|I_t(j)] + \omega_t(j). \tag{40}$$

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

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

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

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

Using the conjectured law of motion for the hierarchy of expectations (34), the expressions for inflation (36) and higher order expectations of marginal cost (42) to write all terms in the Phillips curve (6) as functions of the expectation hierarchy of $\mathbf{x}_t$ we get

$$\mathbf{c}\mathbf{x}_{t|t}^{(0:\infty)} = (1-\theta)(1-\beta\theta)\sum_{k=0}^{\infty}(1-\theta)^k(\mathbf{d}\mathbf{x}_{t|t}^{(k:\infty)} + \mathbf{1}_{1\times 2}\mathbf{x}_{t|t}^{(k)}) + \beta\theta\sum_{k=0}^{\infty}(1-\theta)^k\mathbf{c}M\mathbf{x}_{t|t}^{(k+1:\infty)}, \tag{43}$$

[12] The Kalman filter plays a similar dual role in Woodford (2002).

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

$$\mathbf{c} = \mathbf{a}D + \mathbf{b}\begin{bmatrix} 0 & \mathbf{c}M & & \\ 0 & 0 & \mathbf{c}M & \\ \vdots & & \ddots & \mathbf{c}M \\ 0 & \cdots & \cdots & 0 \end{bmatrix}, \tag{44}$$

where

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

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

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

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

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 we in practise 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 (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 finite number of orders of expectations in the state of the model.[13]

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 Eqs. (38) and (44) and the expression for M and N in the Appendix. 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 our solution, we can simulate the model using (34), (36) and (37).

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.

[13] 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 kth 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 $var(e_t^{(k)}) + var(x_{t|t}^{(k)}) = var(x_{t|t}^{(k-1)})$ and that variances are non-negative yields the desired result.

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 rate 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-rule, $\gamma, \varphi, \phi_\pi$ 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 cost. 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_ε^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 firm's confounding persistent and transitory economy-wide shocks insignificant, but large enough not to make lagged inflation completely revealing of the persistent labour supply shock component ξ_t.

4.1. *Inflation dynamics and the size of the idiosyncratic shocks*

Fig. 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 labour supply shock ξ_t. The solid curve is the response with the idiosyncratic marginal cost shock variance set to zero. The dashed and dotted curves are, respectively, the impulse responses with the idiosyncratic shock variance set to $\frac{1}{2}$ and 2.

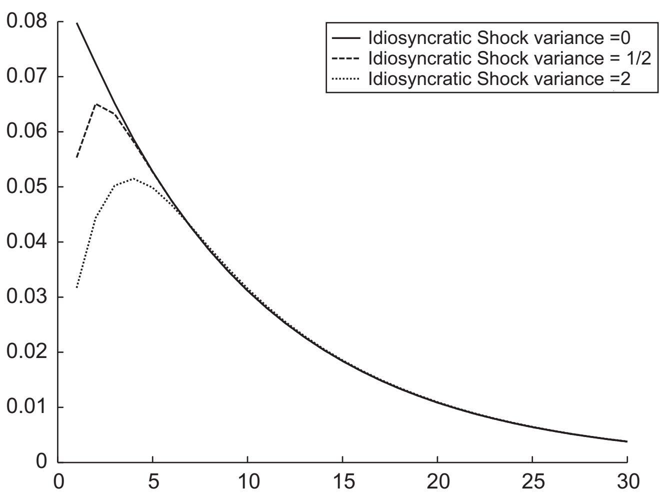

Fig. 1. Impulse response of inflation to unit average marginal cost shock $\sigma_\varepsilon^2/\sigma_v^2 = \{0, \frac{1}{2}, 2\} = \{solid, dashed, dotted\}$.

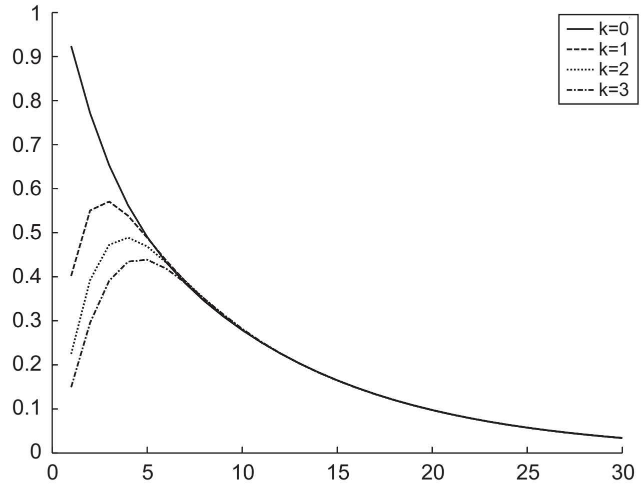

Fig. 2. Impulse response of marginal cost hierarchy to unit average marginal cost shock, $\sigma_\varepsilon^2 = 2$.

There are three things that 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 the variance σ_ε^2 is. Third, the larger this ratio is, the smaller is the first period impact of a marginal cost shock and the lower is the 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 marginal cost. Fig. 2 displays the dynamics of the hierarchy of marginal cost expectations up to the third order, i.e. $\widehat{mc}_t^{(0:3)}$, after a one unit shock to the labour supply shock ξ_t with the variance of the idiosyncratic marginal cost shock set to $\sigma_\varepsilon^2 = 2$.

Fig. 2 shows that the average first order expectation move less than zero order marginal cost on impact. The idiosyncratic component thus works as 'noise' 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 (i.e. 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 impact period response of marginal cost 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'.[14] Herding slows down the social learning process by making the endogenous signal, i.e. 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.

[14]See Morris and Shin (2002) and Chamley (2004).

Table 1
Simulated data and estimates of the Hybrid New-Keynesian Phillips Curve

$\sigma_\varepsilon^2/\sigma_{mc}^2$	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

4.2. The model and U.S. and Euro area inflation dynamics

In the preceding sections we have presented qualitative evidence in the form of hump shaped impulse responses on how imperfect common knowledge can introduce inflation inertia. In this section we ask the question of whether our model can account for the observed inflation inertia in U.S. and Euro area data, with quantitatively realistic amounts of information imperfections. We pursue this question by generating data from our simple general equilibrium model and estimate the Hybrid New-Keynesian Phillips Curve

$$\pi_t = \mu_f \mathrm{E}_t\pi_{t+1} + \kappa mc_t + \mu_b\pi_{t-1} \tag{48}$$

by GMM. If our model is the true data generating process, then the Hybrid New-Keynesian Phillips Curve is of course misspecified. The experiment we perform here is thus to check whether the misspecified econometric model (48) applied to our theoretical model would produce results similar to those obtained when the Hybrid New-Keynesian Phillips Curve is estimated on actual data. Galí et al. (2003) provide a range of estimates for the U.S. 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 U.S. A robust feature across methodologies is that the estimated inflation inertia is lower in the Euro area than in the U.S.

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.[15]

The simulated data was transformed to 'quarterly' frequencies by taking three period averages and the orthogonality condition

$$\mathrm{E}_t[\pi_t - (1-\gamma_b)\pi_{t+1} - \gamma_b\pi_{t-1} - \kappa mc_t] = 0 \tag{49}$$

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

Table 1 tells us that the variance of the idiosyncratic bargaining shock only need to be about $\frac{1}{2}$ the size of the variance of the average marginal cost for the model to generate the observed U.S. inertia. 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% of the firm average wage variations 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 for other industries and countries. However, that the model requires a variance ratio of $\frac{1}{2}$ to match U.S. 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 U.S. consumer price data and that carried out by the

[15] Three hundred 'monthly' observations was transformed into 100 'quarterly' observations. The estimates displayed in the table are averages over 20 independent samples.

Inflation Persistence Network of the central banks within the Euro system. 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 U.S. than for Europe. Bils and Klenow report a median probability of a good changing price in a given month of around 25%. 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 U.S. is around 7 months (rather than (1/.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 U.S. is between 8 and 11 months for the sample period 1998–2005. Alvarez et al find a median duration of consumer prices in Europe of about 10 months.

The microevidence also suggests that individual prices are much more volatile than the aggregate price level. Klenow and Kryvtsov (2004) find that conditional on a price change occurring, the average absolute individual price change is 8.5% which is the same number reported by Nakamura and Steinsson (2007). This can be compared with the average absolute monthly change in the CPI of 0.32% 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 our model, which together with the Calvo mechanism prevents us from analytically deriving the absolute average size of individual price changes. What we can do 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:

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

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

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

since $p_t = \theta p_{t-1} + (1-\theta)\int p_t^*(j)\,\mathrm{d}j$. It is clear from (51) that increased price stickiness, i.e. 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 the variance of the idiosyncratic component is, 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 aggregation, it is only the second indirect effect that affects the aggregate price level. The magnitude of individual relative to aggregate price changes therefore unambiguously increases as the variance of the idiosyncratic component is increased.

In Fig. 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 set to $\frac{1}{2}$ 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 (dotted line) are much more volatile than the aggregate price level (solid line).

We can simulate the aggregate price level path and the path of an individual good's price for various ratios of idiosyncratic and economy-wide marginal cost shock variances and compute the relative size of the absolute price changes of an individual good compared to the absolute average price level changes. The result of this exercise for different degrees of price stickiness is 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 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

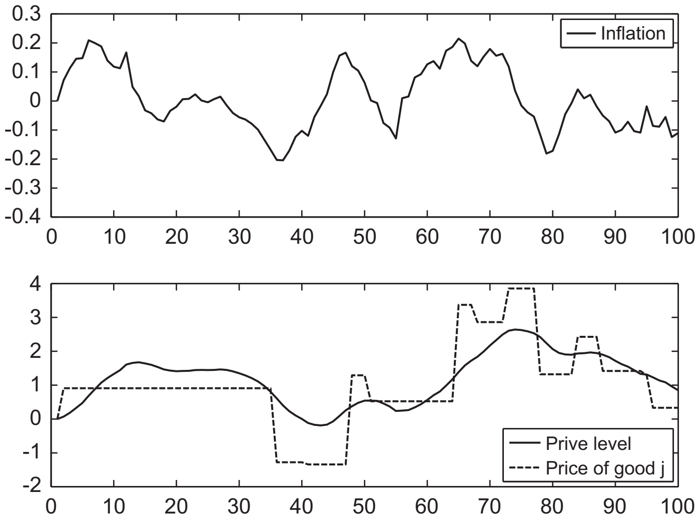

Fig. 3. Inflation, price level and example individual price path.

Table 2
Magnitude of individual and aggregate price changes

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

numbers reported by Klenow and Kryvtsov (2004) 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 we have argued that when firms have idiosyncratic components in their marginal cost, 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 on quantities relevant for 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 in 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 economy-wide components 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

(2005). 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 ours in their more abstract approach. Mackowiak and Wiederholt 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 our model, the precision of firms' information is determined by the relative variance of individual firms marginal cost and the economy-wide aggregate marginal cost. This explicit approach allows us to, at least in principle, compare variances of the model with variances in the data and 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 cost and lagged inflation and output are the only information available to agents in reality. However, a *necessary* condition for limited information availability (or limited capacity to process information) based explanations of economic phenomena to be plausible is that quantities that are immediately and costlessly observable to agents are not too informative. In the specific case considered here, we found that the firm level idiosyncratic wage variances necessary to replicate U.S. inflation dynamics is about $\frac{1}{2}$ of the overall variance of real wages. In principle, we could compare this with variance ratios in the data, but in practice we are constrained by the limited availability of firm level data. Martins (2003) is a rare study that does provide 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 our 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 too strong conclusions about what we should consider to be a realistic lower bound on how imperfectly informed firms can be.

The more explicit approach to information imperfections also points out directions for further research. The present model suggests an explanation for the observed higher inflation inertia in the U.S. relative to the Euro area. European wage bargaining is often centralised, while in the U.S. a larger fraction of wages are set at the firm level.[16] There is thus likely to be more firm level variation in wages in the U.S. than in Europe, which in our model would lead to more inertia, and could thus explain the observed differences. Comparing the predictions of our model with a larger cross section of countries may be another way to validate the main mechanism of the model.

Through the Calvo mechanism the model can be made consistent with the microevidence on the average duration of prices, but it also means that firms need to be forward looking when they set prices. The fact that the agents in the model make dynamic choices, rather than a series of static choices, renders existing solution techniques inapplicable. We solve the model 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, allows us to derive a solution algorithm of arbitrary accuracy. Though some of the details of the algorithm are relegated to the working paper version of this article, it may be of independent interest to some readers.

References

Adam, K., 2007. Optimal monetary policy with imperfect common knowledge. Journal of Monetary Economics 54 (2), 276–301.

Alvarez, L.J., Dhyne, E., Hoeberichts, M.M., Kwapil, C., Le Bihan, H., Lunneman, P., Martins, F., Sabbatini, R., Stahl, H., Vermeulen, P., Vilmunen, J., 2005. Sticky price in the Euro area: a summary of new micro evidence. ECB Working Paper Series No. 563.

Amato, J., Shin, H.S., 2006. Public and private information in monetary policy models. Economic Theory 27, 213–224.

Aucremanne, L., Dhyne, E., 2004. How frequently do prices change? Evidence based on the micro data underlying the Belgian CPI. ECB Working Paper 331.

Bils, M., Klenow, P.J., 2004. Some evidence on the importance of sticky prices. Journal of Political Economy 112, 947–985.

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

Chamley, C.P., 2004. Rational Herds: Economic Models of Social Learning. Cambridge University Press, Cambridge.

[16]See OECD (1997).

Christiano, L.J., Eichenbaum, M., Evans, C., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113, 1–45.

Fuhrer, J., Moore, G., 1995. Inflation persistence. Quarterly Journal of Economics 110, 127–160.

Gali, J., Gertler, M., 1999. Inflation dynamics: a structural econometric analysis. Journal of Monetary Economics 44, 195–222.

Galí, J., Gertler, M., López-Salido, J.D., 2003. Erratum to ''European Inflation Dynamics''. European Economic Review 47, 759–760.

Hellwig, C., 2004. The social cost of heterogenous information. Mimeo UCLA.

Klenow, P.J., Kryvtsov, O., 2004. State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? NBER Working Paper 11043.

Lorenzoni, G., 2007. News shocks and optimal monetary policy. NBER Working Paper 12898.

Lucas Jr., R.E., 1972. Expectations and the neutrality of money. Journal of Economic Theory 4, 103–124.

Lucas Jr., R.E., 1973. Some international evidence on output-inflation trade offs. American Economic Review 63, 326–334.

Lucas Jr., R.E., 1975. An equilibrium model of the business cycle. Journal of Political Economy 83, 1113–1144.

Mackowiak, B., Wiederholt, M., 2005. Optimal sticky prices under rational inattention. Mimeo, Humboldt University.

Mankiw, G., Reis, R., 2002. Sticky Information vs. sticky prices: a proposal to replace the New Keynesian Phillips curve. Quarterly Journal of Economics 117, 1295–1328.

Martins, P.S., 2003. Firm wage differentials in a competitive industry: some matched panel evidence. International Journal of Manpower 24, 336–346.

Morris, S., Shin, H.S., 2002. Social value of public information. American Economic Review 92, 1521–1534.

Nakamura, E., Steinsson, J., 2007. Five facts about prices: a reevaluation of menu cost models. Mimeo, Harvard University.

OECD, 1997. Economic performance and the structure of collective bargaining. In: OECD Employment Outlook, July 1997. OECD, Paris.

Phelps, E.S., 1970. Introduction: the new microeconomics in employment and inflation theory. In: Phelps, E.S., et al. (Eds.), Microeconomic Foundations of Employment and Inflation Theory. Norton, New York.

Sims, C.A., 2003. Implications of rational inattention. Journal of Monetary Economics 50, 665–690.

Smets, F., Wouters, R., 2002. An estimated dynamic stochastic general equilibrium model for the Euro area. Journal of the European Economic Association 1, 1123–1175.

Taylor, J.B., 1993. Discretion versus policy rules in practice. Carnegie-Rochester Conference Series on Public Policy 39, 195–214.

Woodford, M., 2002. Imperfect common knowledge and the effects of monetary policy. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (Eds.), Knowledge, Information, and Expectations in Modern Macroeconomics: In Honour of Edmund S. Phelps. Princeton University Press, Princeton.

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