Setting the Right Prices for the Wrong Reasons

Christian Hellwig[*] — UCLA and CEPR

Venky Venkateswaran[†] — UCLA

December 2008

Abstract

We consider a model of nominal price adjustment with firm-specific and aggregate shocks to economic fundamentals and incomplete, dispersed information. Firms update their expectations about fundamentals based on their own cash flows (revenues and wages). We show that a model with realistic levels of product-level price dispersion, the firms' inference about aggregate shocks is very gradual, yet in the aggregate prices adjust rapidly in response to aggregate nominal shocks. When an aggregate shock occurs, firms mistakenly attribute it to firm-specific shocks, but adjust prices nevertheless, since the exact nature of the shock matters little for its optimal pricing decision.

1 Introduction

Understanding the relation between monetary policy, monetary aggregates, prices and income at various horizons is a central issue in monetary economics. One potential explanation of monetary non-neutrality at short- to medium-run horizons, going back to Phelps (1970) and Lucas (1972), is based on lack of information. If economic agents don't have real time access to all relevant information, and they face uncertainty about the current economic conditions, their decisions will reflect changing economic fundamentals only gradually, as the information becomes available.

One drawback of the original incomplete information models was the lack of an internal amplification mechanism - the speed of price adjustment and the degree of non-neutrality was directly and exclusively related to the arrival rate of new information. This implies that prices should fully

[*]Email: chris@econ.ucla.edu

[†]Email: vvenkate@ucla.edu

incorporate all new information as it becomes available. In part because of concerns such as these, the monetary policy literature has by and large abandoned incomplete information as a source of non-neutrality, and instead focused on rigidities and adjustment costs in price-setting.

Recently, the incomplete information hypothesis has seen a revival with the works of Woodford (2001) and Mankiw and Reis (2002), who argued that heterogeneity in beliefs may amplify the effects of incomplete information, when pricing decisions are complementary across firms. Woodford (2001) shows that, with heterogeneous information, firms face uncertainty not only about aggregate fundamentals, but also about the pricing decisions of other firms. When pricing decisions are complementary, firms become a lot more reluctant to incorporate new information, because they can no longer be sure that other firms share the same assessment of fundamentals, and hence will also adjust their prices.[1]

In this paper, we provide a framework of analysis and a quantitative evaluation for this new class of nominal business cycle models with incomplete, heterogeneous information. We formulate a dynamic, stochastic general equilibrium model in which firms have limited information about the underlying economic shocks. The model orients itself along the lines of canonical New Keynesian Sticky Price Models a la Woodford (2003), but replaces the assumption that prices are exogenously sticky with an assumption that firms don't have access to complete information about all underlying economic shocks. We enrich the model to allow for firm-specific as well as aggregate shocks to economic fundamentals, and we postulate that firms update their expectations based on the information generated by their market activities, i.e. firm-level revenues and costs, which are impacted by the aggregate as well as firm-specific fundamentals.

We use this model to explore analytically and quantitatively the effects of heterogeneous information for price adjustment in response to aggregate nominal shocks. We calibrate our model to match observations about price fluctuations at the product level, in particular the observed level of cross-sectional price dispersion, and the large transitory price movements. As our main conclusion from this calibration, we argue that in contrast to the recent literature following Woodford (2001), heterogeneous information is unable to generate any quantitatively meaningful short-run real effects from nominal shocks, despite the fact that firms have access to only a very limited

[1]While multiple authors have followed up on this with different interpretations for the sources of disagreement (e.g. Inattentiveness or infrequent, staggered information updating in Mankiw and Reis, (2002), rational inattention or information processing constraints in Sims, (2003) and Woodford, (2001), and Mackowiak and Wiederholt, (2008), or spacial separation in an island structure as in Lorenzoni (2008)), the basic insight that disagreement in expectations increases persistence, when coupled with decision complementarities, is robust across these models. See Hellwig (2008b) for an overview of the recent literature on incomplete information business cycle models.

amount of information, and pricing decisions are highly complementary.

To understand the reason underlying this stark result, one must turn to the additional features of our model, in particular, the presence of large idiosyncratic shocks that are transient but not completely transitory. These shocks complicate the inference problem, and one might imagine that, since idiosyncratic shocks are an order of magnitude larger than aggregate shocks, this must render inference of aggregate conditions nearly impossible, and hence significantly slow down aggregate price adjustment.

However, the idiosyncratic shocks also provide firms with additional motives for adjusting prices, which in turn makes the inference problem less relevant for firm pricing, and mitigates the amplifying effects of information heterogeneity and pricing complementarities. Firms care about separating one type of shock from the other only to the extent that their pricing responses to the two shocks are different. For the firm, however, it makes little difference whether its costs rise because of idiosyncratic or aggregate shocks, as long as its response of increasing its price remains approximately the same.

At this point, a second important assumption comes into play; namely, we are assuming that firms extract information from their own market transactions. The resulting market-generated demand and cost signals turn out to be a parsimonious, yet very accurate indicator of the firms' optimal pricing decisions in real time, despite being a poor source of information for learning about aggregate vs. firm-specific shocks. As a result, when an aggregate shock occurs, the firms are confused about its nature and will mistakenly attribute it to idiosyncratic fluctuations, but they nevertheless will adjust prices, because the exact nature of the shock is less relevant for the pricing decisions. In other words, prices adjust quickly, but for the 'wrong' reasons.

That said, we do not obtain complete convergence of prices in response to aggregate shocks in the short- to medium run. After a large initial adjustment to 'almost' the right level for the reason explained above, it takes firms a very long time to move from almost complete to complete price adjustment; that is, although the model generates only small real effects, these tend to be very highly persistent. Complete price adjustment requires firms to completely separate idiosyncratic from aggregate shocks, and wiht idiosyncratic shocks being an order of magnitude larger than aggregate ones, the learning about aggregate shocks can take a very long time. Obviously this implication is not robust to the inclusion of additional sources of information.[2]

Our paper is related to several branches of the literature. For our calibration, we draw on the

[2]By allowing for only a very limited set of information sources, we have given the model the best possible chance of generating large adjustment delays from information heterogeneity.

large recent literature documenting price adjustment at the micro level using large scale data sets of individual price quotes. Three observations about product-level fluctuations are particularly relevant for our quantitative conclusions: First, we draw on Klenow and Kryvtsov (2008) and Burstein and Hellwig (2007) for measures of price dispersion and the large idiosyncratic fluctuations in prices. These fluctuations are transient, but not completely transitory. Moreover, they are too large to be consistent with a model that, like Woodford (2001), focuses on information as the only source of heterogeneity, which suggests that firm-specific fundamental shocks must play an important role for product-level price adjustment. Second, we draw on Midrigan (2007) for a measure of serial correlation in prices at a monthly frequency. This measure in turn is consistent with the serial correlation typically assumed in calibration of menu cost models to micro data and plays an important role in pinning down the importance of the idiosyncratic pricing motive. Finally, we draw on Burstein and Hellwig (2007) and Eichenbaum, Jaimovich and Rebelo (2008) for observations about quantity fluctuations and correlations between prices and quantities. These authors both document that prices and quantities are at best weakly negatively correlated, and they often move in the same direction, suggesting that prices respond to cost as well as demand fluctuations. This in turn suggests that both cost and demand shocks must be sufficiently persistent to generate large pricing responses.

Our conclusion of a high degree of price flexibility is in stark contrast to the results obtained by Mackowiak and Wiederholt (2008), and it is therefore worthwhile discussing the differences behind these results. Mackowiak and Wiederholt (2008) study a model of price adjustment with rational inattention, in which firms have to divide a limited information processing capacity between tending to idiosyncratic technology and aggregate demand shocks. Since idiosyncratic technology shocks are an order of magnitude larger than aggregate demand shocks, firms tend to pay very little attention to the latter, so that nominal shocks can have persistent real effects. As a result, prices are sticky in response to aggregate demand shocks, yet very responsive to firm-specific technology shocks.

The key distinction between our model and theirs is that we do not allow for a complete separation between learning about cost- and demand-side conditions. In particular, we can replicate their findings in our model by (i) making the firm-specific wage and demand shocks completely transitory, so that they have no effect on firms' optimal prices and serve purely as informational noise, and (ii) adding idiosyncratic technology shocks with a magnitude and persistence that matches the observed price dispersion. However, such a calibration has the counter-factual implication that prices and quantities are highly negatively correlated.

On the methodological side, our paper contributes to the formulation and solution of general equilibrium models with heterogeneous information. Using techniques first suggested in Hellwig (2002) and fully developed in Hellwig (2008a), we propose a simple and computationally efficient algorithm for solving our model. A key problem in heterogeneous information models is that the model does not easily admit a recursive equilibrium characterization on a finite-dimensional state space. In equilibrium, firms need to make forecasts about the forecasts and pricing decisions of other firms, which in turn depend on their forecasts of others, and so on... (see Townsend (1983)). This problem of forecasting the forecasts of others results because of (i) the fact that the firms' best response pricing strategies depend on their forecasts of what other firms are doing, and (ii) firms will then need to form forecasts of the other firm strategies in order to interpret the market-generated information, if the latter is endogenously generated from market transactions, and thus dependent on equilibrium pricing strategies. Both of these reasons render the equilibrium filtering problem highly intractable, and typically deprive the model of finite-dimensional recursive equilibrium structure.

Instead of looking for a recursive solution to the model, we simply assume that shocks become common knowledge after a finite, arbitrarily large delay. This lets us recast the equilibrium characterization as a static, finite-dimensional signal extraction problem, which turns out to have a very tractable structure, with an easily interpretable solution that is almost in closed form and easy to implement computationally. Although we develop the solution for a specific model, the method can be generalized to almost arbitrary information structures and requires little to no assumptions about the economic model other than a log-linearization around a steady-state.

In section 2, we show by means of a simple example, how the addition of idiosyncratic reasons for price changes alters aggregate implications of heterogeneous information and price adjustment, mitigating in particular the effects of pricing complementarity. The example also illustrates the tension that arises within the model between generating substantial amounts of price dispersion, and generating large aggregate output effects from nominal shocks. In section 3, we formulate our general model, and develop methods for equilibrium characterization, along with some analytical results. In section 4, we discuss our calibration strategy and the quantitative implications of our model. Section 5 presents extensions of our benchmark model, which are followed by a brief conclusion.

2 A simple example

We illustrate our main insights using a stylized example. Consider a measure-1 continuum of firms, whose optimal pricing decisions are characterized by the following best-response relation:

$$p_i = \mathbb{E}_i\left(p_i^*\right) + r\mathbb{E}_i\left(p - p^*\right) \tag{1}$$

where p_i denotes firm i's optimal price, $p = \int p_i di$ denotes the average price of all firms, p_i^* and $p^* = \int p_i^* di$ denote exogenous, stochastic 'target' levels for firm i's price and for the aggregate price level, respectively, and $\mathbb{E}_i\left(\cdot\right)$ denotes the firm's expectations conditional on its available information. If the firms had complete information of the targets p_i^* and p^*, there would be full adjustment of all prices to the target levels in equilibrium: $p_i = p_i^*$ and $p = p^*$; p_i^* and p^* can thus be interpreted as the full information equilibrium levels. The parameter $r \in (-1, 1)$ measures the degree of complementarity or substitutability in the firms' pricing decisions, i.e. the extent to which an individual firm's optimal price is increasing or decreasing in the average price of the other firms.

The targets are stochastic, and not directly observable. Let us suppose that firm i's target is given by

$$p_i^* = \gamma z_i + m, \text{ and hence } p^* = m,$$

where $m \sim \mathcal{N}\left(0, \sigma_m^2\right)$ denotes an aggregate shock to the target level, and $z_i \sim \mathcal{N}\left(0, \sigma_z^2\right)$ denotes an idiosyncratic shock. Substituting into the best response function (1),

$$p_i = \gamma\mathbb{E}_i[z_i] + \mathbb{E}_i[m] + r\mathbb{E}_i[p - m] \tag{2}$$

The firm in turn observes a signal s_i that is a linear combination of these two shocks:

$$s_i = z_i + m + \zeta_i,$$

where $\zeta_i \sim \mathcal{N}\left(0, \sigma_\zeta^2\right)$ represents idiosyncratic signal noise. $\mathbb{E}_i\left(\cdot\right) = \mathbb{E}\left(\cdot|s_i\right)$ then denotes the firm's expectation conditional on its signal s_i. The interpretation of this information structure, which we will render explicit in our full-fledged model, is the following: the firms face idiosyncratic and aggregate fluctuations in their demand, cost or productivity levels. They extract information from the information that is generated by their market activities (such as their sales, their productivity level, or their wage bill). This information reflects fluctuations in both idiosyncratic and aggregate conditions, but does not enable the firms to fully isolate what fluctuations are due to idiosyncratic shocks, and what is due to aggregate shocks. The extent to which firms may want to respond to one type of shock but not the other is captured by the parameter γ.

The example also incorporates a number of notable special cases. If $\gamma = 0$, the idiosyncratic shock does not affect payoffs, and the distinction between z_i and ζ_i loses its meaning. In this case, the only source of heterogeneity is informational - loosely speaking, this corresponds to the class of heterogeneous information models that were initiated by Woodford (2001). If on the other hand, $\gamma = 1$, we have a special case where the agents observe a signal of their target price level. Moreover, when $\sigma_\zeta^2 = 0$, this signal becomes perfect. In that case, each firm would be able to set its price exactly equal to its target, and it will be willing to do so, if it can count on all other firms doing the same.

Equilibrium Characterization: To characterize equilibria in this simple pricing game, we conjecture a linear pricing rule of the form $p_i = ks_i$. Then, the average price satisfies $p = kz$, and substituting into the firms' best response function (1), we find:

$$p_i = \mathbb{E}_i\left(p_i^*\right) + r\mathbb{E}_i\left(p - p^*\right) = \mathbb{E}_i\left(\gamma z_i + m - r\left(1-k\right)m\right).$$

By Bayesian updating, $\mathbb{E}_i\left(m\right) = \sigma_m^2 / \left(\sigma_m^2 + \sigma_z^2 + \sigma_\zeta^2\right) \cdot s_i$ and $\mathbb{E}_i\left(z_i\right) = \sigma_z^2 / \left(\sigma_m^2 + \sigma_z^2 + \sigma_\zeta^2\right) \cdot s_i$. Substituting this into the previous equation, we verify the conjecture and solve for k:

$$k = \frac{\left(1-r\right)\sigma_m^2 + \gamma\sigma_z^2}{\left(1-r\right)\sigma_m^2 + \sigma_z^2 + \sigma_\zeta^2}. \tag{3}$$

This parameter k thus summarizes both the response of individual firms' prices to their signals, and the response of the aggregate price level to aggregate shocks.

Implications for aggregate price adjustment: With this characterization, we can note a number of results, both old and new. We begin with the case when $\gamma = 0$, in which case the only source of heterogeneity is informational, and z_i can be interpreted as a purely informational shock.

Proposition 1 *When $\gamma = 0$, a higher degree of complementarity ($\frac{\partial k}{\partial r} < 0$), or a lower precision of the idiosyncratic shock reduce the response of prices to aggregate shocks.*

This proposition summarizes the idea that in the presence of heterogeneous information, pricing complementarities amplify the adjustment delay of prices to aggregate shocks: when $r = 0$, $k = \sigma_m^2 / \left(\sigma_m^2 + \sigma_z^2 + \sigma_\zeta^2\right)$, so prices respond to the aggregate shocks with a rate that just reflects the signal to noise ratio $\sigma_m^2 / \left(\sigma_z^2 + \sigma_\zeta^2\right)$. When instead pricing decisions are complementary ($r > 0$), strategies discount the private signals relative to their information content in a way that is directly proportional to the complementarity. In this case, firms attempt to forecast not only the common target price level, but also the other firms' forecasts of the target. While the former relies on the firms' private signals just in proportion to their information content, the firms higher-order forecasts of the other firms expectations puts higher weights on the firm's priors or other common sources of information, as these are more directly informative of the other firms' beliefs. In equilibrium, the firms' optimal pricing response thus discounts the information contained in the private signals, and amplifies the weight attributed to the prior, as the latter is disproportionally more useful in forecasting the other firms' prices. As $r \to 1$, this discount can become arbitrarily large, and hence in the limit, firms may completely discount private signals, even if they are very precise, because they do not expect other firms to respond to their private information.

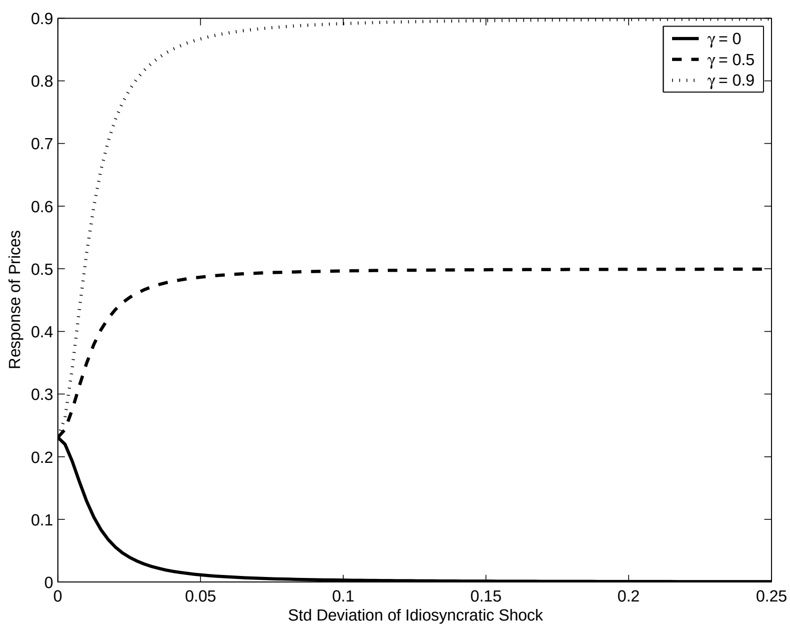

Figure 1: Effect of σ_z^2 on Price Response

The next proposition summarizes how this finding is affected when the firms' optimal pricing decisions also depend on the idiosyncratic shocks z_i:

Proposition 2 *(i) An increase in γ increases the response of prices to aggregate shocks ($\frac{\partial k}{\partial \gamma} > 0$). The response of prices to aggregate shocks is bounded below by $R \equiv \gamma \cdot \sigma_z^2 / \left(\sigma_z^2 + \sigma_\zeta^2\right) > 0$.*

(ii) When $R < 1$, pricing complementarities still delay adjustment ($\frac{\partial k}{\partial r} < 0$), but an increase in γ mitigates their effects ($\frac{\partial^2 k}{\partial \gamma \partial r} > 0$).

(iii) When $R = 1$, the equilibrium generates full adjustment to aggregate shocks i.e. $k = 1$. Complementarities and Heterogeneity in information have no effect on the equilibrium. Moreover, if $\sigma_\zeta^2 = 0$ (and

hence $\gamma = 1$), the equilibrium exactly replicates the full information equilibrium.

(iv) When $R > 1$, $k > 1$, and the effects of pricing complementarities are overturned: complementarities amplify the response to aggregate shocks ($\frac{\partial k}{\partial r} > 0$), and their effects are reinforced if γ increases.

The addition of idiosyncratic shocks to the firm's pricing target thus significantly alters the conclusions of Proposition 1. The parameter R, i.e. the effect of idiosyncratic shocks on optimal prices multiplied with the signal-to-noise ratio for the idiosyncratic shock, serves as a lower bound on the response to aggregate shocks - we can thus no longer obtain arbitrary amplification effects from complementarities. When $R < 1$, the effects of pricing complementarities still hold, but are weakened by the fundamental shocks. To understand these results, suppose first that the firms thought that σ_m^2 was zero, i.e. that they never faced any aggregate shocks, and suppose now that there was a one-time increase in z, which appears as an increase in the firms' signals. The firms would mistakenly attribute this shock to a combination of the idiosyncratic state and noise, and would respond to this shock with a rate R, i.e. their optimal full information rate adjusted for the signal noise. Now, if the variance of an aggregate shock was not zero, the firms' inference and optimal pricing decisions will also allow for that, but this matters only for the response in excess of R, and it is only on this remainder that the pricing complementarity has its effect of delaying adjustment.

When $R = 1$, the adjustment resulting from an efficient response to idiosyncratic shocks then implies a full adjustment to the aggregate shock. In this case, firms maybe highly uncertain about whether changes in the signal result from idiosyncratic or aggregate shocks, but this uncertainty doesn't affect their decisions. For example, when $\gamma = 1$ and $\sigma_\zeta^2 = 0$, signals provide perfect information about the current period's target under full information. The resulting equilibrium displays full price adjustment, despite the fact that firms remain imperfectly informed about the magnitudes of idiosyncratic and aggregate shocks, and despite the fact that they face a motive to coordinate pricing decisions. Firms will then set exactly the right prices, albeit for the wrong reasons.

Implications for price dispersion: As the previous discussion made clear, the extent to which idiosyncratic shocks influence the firms' pricing targets is key for the aggregate implications of heterogeneous information. A quantitative assessment of such an information channel thus requires imposing some discipline on the parameter choice of γ and σ_z. To conclude the discussion of this simple example, we explore how the use of micro data, in particular on dispersion of prices at the micro level, can help us draw quantitative conclusions.

In this simple example, price dispersion, measured as the cross-sectional standard deviation of prices at the product level, is given by

$$Disp = \left\{ \int (p_i - p)^2 \, di \right\}^{1/2} = k \cdot \sqrt{\sigma_z^2 + \sigma_\zeta^2} = \left[\frac{(1-r)\,\sigma_m^2 + \gamma \sigma_z^2}{(1-r)\,\sigma_m^2 + \sigma_z^2 + \sigma_\zeta^2} \right] \sqrt{\sigma_z^2 + \sigma_\zeta^2}$$

We begin again with the benchmark case when $\gamma = 0$. In this case, dispersion is increasing for low levels of σ_z^2, and decreasing for high levels, and reaches a maximal level when $(1-r)\,\sigma_m^2 = \sigma_z^2 + \sigma_\zeta^2$. Dispersion is therefore bounded above by

$$Disp \leq \frac{\sqrt{1-r}}{2} \sigma_m.$$

Thus, a model purely based on informational heterogeneity will imply a level of cross-sectional price dispersion that is necessary smaller in magnitude than the standard deviation of the aggregate shock to the target, and for commonly used degrees of pricing complementarity, the gap may be quite substantial (if $r = 0.84$, $Disp \leq 0.2\sigma_m$).[3] In practice, we see the opposite, i.e. cross-sectional price dispersion tends to be an order of magnitude larger than the standard deviation of aggregate shocks.

How is this conclusion affected by γ? When $\gamma > 0$, any level of price dispersion may be the result of sufficiently large idiosyncratic shocks. To see this, notice that $\lim_{\sigma_z^2 \to \infty} Disp = \infty$, for any $\gamma > 0$. Price dispersion is then driven entirely by the fact that the equilibrium response to the signal is bounded away from zero, regardless of whether it is the result of an idiosyncratic or an aggregate shock.

To summarize, our example illustrates that there is a tension for heterogeneous information models between generating significant levels of price dispersion and meaningful delays in aggregate price adjustment. Large delays of price adjustment are most likely to result from heterogeneous information, when complementarities are strong (i.e. r close to 1), and optimal prices do not respond to idiosyncratic shocks (γ is small). However, heterogeneity in information is, by itself, unable to generate quantitative meaningful levels of price dispersion, which points to an important role for additional idiosyncratic payoff shocks, requiring positive γ and positive σ_z^2. This in turn reduces adjustment delays.

[3] What is more, the computation of this bound allowed for the maximal level of dispersion, at which exactly $1/2$ of the aggregate shock gets absorbed by prices. Assuming any stronger degree of adjustment delays would further reduce cross-sectional price dispersion.

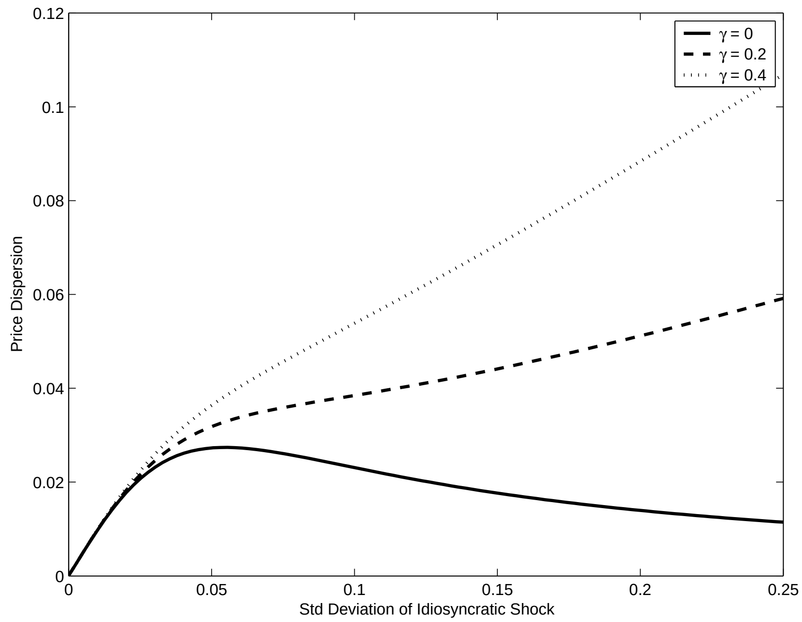

Figure 2: Effect of σ_z^2 on Price Dispersion

A high γ, for example, increases dispersion, but also increases k. Such a calibration, therefore, matches price dispersion at the cost of significant real effects of aggregate shocks. The model may have its best chance to achieve both when γ is low and σ_z^2 is high, so that the idiosyncratic payoffs shocks matter little for the price targets, yet they are big enough to generate substantial price dispersion. In the remainder of this paper, use the additional structure imposed by a fully micro-founded, quantitative, dynamic general equilibrium model to assess the plausibility of such a parameter choice.

3 The general model

In this section, we formulate a fully micro-founded dynamic stochastic general equilibrium model in which information heterogeneity emerges endogenously from the assumption that firms have only limited access to information. Our working assumption is that at a minimum, firms observe the information that is generated through their market transactions, which consist of the demand for their products (or equivalently, the firms' revenues), and of the wages paid to their workers. Wages and demand observations are subject to idiosyncratic and aggregate fluctuations, and do not allow the firms to fully infer the underlying states.

In order to focus on the role of pricing complementarities and information vs. payoff heterogeneity on the production side, we will keep the household side of the model as close to the New Keynesian benchmark as possible. In particular, we assume that the representative household has access to complete information about the aggregate and market-specific shocks, and has access to complete contingent claims markets.

3.1 Model description

Representative Household: The representative household maximizes its lifetime utility over consumption C_t and real balances M_t/P_t, as well as disutility of effort over a measure 1 continuum of labor types N_{it},

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\psi}}{1-\psi} + \ln \frac{M_t}{P_t} - \frac{1}{1+\kappa} \int_0^1 Z_{it} N_{it}^{1+\kappa} di \right),$$

where $\beta \in (0,1)$, $\psi > 0$, $\kappa > 0$, and N_{it} represents labor supply and Z_{it} , an idiosyncratic preference shock for labor of type i. $\mathbb{E}_t(\cdot)$ denotes the representative household's expectations as of date t. Assuming that the household has access to a complete contingent claims market, we can write the household's life-time budget constraint as

$$M_0 \geq \mathbb{E}_0 \sum_t \lambda_t \{ C_t P_t + i_t M_t - \int W_{it} N_{it} di - \Pi_t - T_t \}$$

where λ_t denotes the economy's stochastic discount factor used to price nominal balances, Π_t and T_t denote aggregate corporate profits and taxes or transfers (in nominal terms), W_{it} denotes the nominal wage for labor of type i, and the term $i_t M_t$ denotes the household's opportunity costs of holding monetary balances at date t. The first order conditions of this problem are

$$\begin{aligned}
\beta^t C_t^{-\psi} &= \lambda_t P_t \\
\beta^t \frac{1}{M_t} &= \lambda_t i_t \\
\beta^t \eta Z_{it} N_{it}^{\kappa} &= \lambda_t W_{it} \\
\lambda_t &= (1+i_t) \mathbb{E}_t \lambda_{t+1}
\end{aligned}$$

Along with the budget constraint and a law of motion for i_t that is determined by monetary policy, these equations characterize the solution to the household's problem. Throughout this paper, we focus on the special case where $\ln M_t$ follows a random walk with drift μ:

$$\ln M_{t+1} = \ln M_t + \mu + \sigma_u u_t$$

where u_t is an iid random variable, distributed $N(0,1)$. This implies that interest rates are constant at a level $\hat{\imath}$, defined by $(1+\hat{\imath})^{-1} = \beta \mathbb{E}_t (M_t/M_{t+1}) = \beta exp\left(-\mu + \sigma_u^2/2\right)$, which we assume to be strictly positive. Along with the FOC, we can then write state prices, consumption, and wages as follows:

$$\lambda_t = \beta^t \frac{1}{M_t} \tag{4}$$

$$C_t = K_0 \left(\frac{M_t}{P_t}\right)^{\frac{1}{\psi}} \tag{5}$$

$$W_{it} = K_1 M_t Z_{it} N_{it}^{\kappa} \tag{6}$$

where K_0 and K_1 are time-independent constants. Besides this characterization of the household's equilibrium behavior through static equations for aggregate consumption and type-specific wages, the constant interest rate benchmark also eliminates the effects of nominal interest rates as a public signal of aggregate economic activity.[4]

Production Side: Consumption C_t is a composite good, assembled by competitive firms using a standard Dixit-Stiglitz aggregator

$$C_t = \left(\int B_{it}^{\frac{1}{\theta}} C_{it}^{\frac{\theta-1}{\theta}} di\right)^{\frac{\theta}{\theta-1}}$$

where B_{it} is an idiosyncratic shock process and C_{it} is the amount of output of sector i used in the production of the final good. This leads to the standard expressions for demand functions for the output of sector i

$$C_{it} = B_{it} C_t \left(\frac{P_{it}}{P_t}\right)^{-\theta} \tag{7}$$

and the aggregate price index

$$P_t = \left(\int B_{it} P_{it}^{1-\theta} di\right)^{\frac{1}{1-\theta}}. \tag{8}$$

Each intermediate good i is produced by a single firm, using labor type i as its unique input into production. We assume that production in intermediate sector i is linear in the labor input, and is given by

$$Y_{it} = A_{it} N_{it} \tag{9}$$

For the moment[5], we assume that productivity is constant at 1 i.e. $A_{it} = 1$.

The intermediate producer in sector i sets a price P_{it} to solve

$$\max_{P_{it}} \mathbb{E}_{it} \left(\lambda_t \left(P_{it} C_{it} - N_{it} W_{it}\right)\right).$$

[4]It is conceptually straight-forward to extend our analysis to a more complete monetary policy framework. Moreover, the additional information conveyed by interest rates would reduce both the level of price dispersion and the degree of non-neutrality, which we conjecture would reinforce the results presented here.

[5]In section 5, we will consider extensions with idiosyncratic and aggregate productivity shocks.

These prices are set before markets open, and the expectation $\mathbb{E}_{it}(\cdot)$ is conditional on the information available to the firm at the time of making its pricing decision. The firm's information set $\mathcal{I}_t^i$ will be defined below.

Substituting from the optimality conditions of the household (4)-(6) and the demand function (7), the problem can be written as

$$\max_{P_{it}} \mathbb{E}_{it}\left[B_{it}\frac{C_t P_t}{M_t \hat{\imath}}\left(\frac{P_{it}}{P_t}\right)^{1-\theta} - K_1 Z_{it}\left(B_{it}C_t\left(\frac{P_{it}}{P_t}\right)^{-\theta}\right)^{1+\kappa}\right]$$

The resulting first order condition is

$$P_{it}^{1+\theta\kappa} = \frac{\theta(1+\kappa)\hat{\imath}}{\theta-1}\cdot\frac{\mathbb{E}_{it}[Z_{it}B_{it}^{1+\kappa}C_t^{1+\kappa}P_t^{\theta(1+\kappa)}]}{\mathbb{E}_{it}[B_{it}C_t P_t^{\theta}M_t^{-1}]}$$

Under the assumption that (i) conditional on the firm's information, $Z_{it}b_{it}^{1+\kappa}C_t^{1+\kappa}P_t^{\theta(1+\kappa)}$ and $B_{it}C_tP_t^{\theta}M_t^{-1}$ are log-normally distributed, and (ii) conditional on a realization of aggregate shocks, the cross-sectional distribution of prices across firms is also log-normal, we can take logs on both sides and find:[6]

$$p_{it} = Const + \frac{\kappa}{1+\theta\kappa}\mathbb{E}_{it}[b_{it}+c_t+\theta p_t] + \frac{1}{1+\theta\kappa}\mathbb{E}_{it}[z_{it}+m_t] \tag{10}$$

for some real-valued constant $Const$. Using the household FOC (5) to substitute for $\ln C_t$, this becomes

$$p_{it} = Const + \frac{1}{1+\theta\kappa}\mathbb{E}_{it}[\kappa b_{it}+z_{it}] + \mathbb{E}_{it}[m_t] + \frac{\kappa\left(\theta-\psi^{-1}\right)}{1+\theta\kappa}\mathbb{E}_{it}[p_t-m_t], \tag{11}$$

$$\text{where } p_t = \ln P_t = \frac{1}{1-\theta}\ln\left(\int_0^1 B_{it}P_{it}^{1-\theta}di\right) = const + \int_0^1 p_{it}di.$$

For a given information structure $\left\{\mathcal{I}_t^i, i\in[0,1]\right\}$, a price function $p\left(\mathcal{I}_t^i\right)$ then characterizes an *equilibrium*, if and only if it constitutes a fixed point of the firms' optimal pricing condition (11). Notice that (11) represents exactly the form of equilibrium pricing relation (1) that we used in the simple example of the previous section. In particular, $r \equiv \frac{\kappa\left(\theta-\psi^{-1}\right)}{1+\theta\kappa}$ determines the degree of pricing complementarities (corresponding to r in the previous section's notation). The firms' willingness to respond to idiosyncratic shocks is determined in part by $\kappa/(1+\theta\kappa)$ and $1/(1+\theta\kappa)$, for demand and cost shocks, respectively.

Information structure, idiosyncratic and aggregate shocks: To complete the model description, we describe the information structure $\left\{\mathcal{I}_t^i, i\in[0,1]\right\}$ and specify the stochastic processes for

[6] We shall use small letters to denote the natural logs of capital-lettered variables, e.g. for any variable X, we write $x = \ln X$.

the various shocks. The model has a single aggregate shock to supply m_t, which follows a random walk with drift μ and variance σ_u^2, as well as product-specific shocks to demand B_{it} and labor supply Z_{it}. We assume that the two product-specific shocks are each the sum of a persistent component that admits an AR(1) representation, and a transitory shock that is iid over time. That is, B_{it} and Z_{it} admit the following representations:

$$b_{it} = \sigma_b \sum_{s=0}^{\infty} \rho_b^s v_{i,t-s}^1 + \tilde{\sigma}_b \tilde{v}_t^1 \qquad\qquad z_{it} = \sigma_z \sum_{s=0}^{\infty} \rho_z^s v_{i,t-s}^2 + \tilde{\sigma}_z \tilde{v}_t^2 \tag{12}$$

where v_t^1, $\tilde{v}_t^1$, v_t^2, and $\tilde{v}_t^2$ are iid random variables distributed according to $\mathcal{N}(0,1)$, and ρ_b, σ_b, $\tilde{\sigma}_b$, ρ_z, σ_z and $\tilde{\sigma}_z$ are all non-negative.

In each period, a firm observes the demand for its product C_{it}, as well as its wage bill $W_{it}N_{it}$. This is informationally equivalent to observing one demand signal x_{it} and a wage signal ω_{it}. Using the demand function (7) and the household optimality conditions (4)-(6), these signals can be written as a function of the exogenous shocks and the aggregate price level:

$$x_{it} = b_{it} + \psi^{-1} m_t + (\theta - \psi^{-1}) p_t \tag{13}$$

$$\omega_{it} = z_{it} + m_t \tag{14}$$

The idiosyncratic demand and wage shocks thus prevent the full revelation of aggregate shocks from demand and cost observations. Nevertheless, firms are able to infer their target prices within each period from these signals - in fact, if the firms had the information about the current signal realization in advance of setting their current prices, then by setting $p_{it} = \kappa/(1+\theta\kappa) \cdot x_{it} + 1/(1+\theta\kappa) \cdot \omega_{it}$, they would be able to set exactly the price that is optimal for the current period.[7]

In our model, this is prevented by the assumption that a firm has to decide on its price in advance of the realization of current market outcomes and current wage and demand signals. Its expectations are then based on realizations of the signals up until the preceding period. This timing assumption enables us to link the specification of idiosyncratic shocks back to our stylized example. The transitory components will play a similar role as the signal noise, while the persistent component will play the role of the firm-specific shocks to price targets, with the degree of persistence determining the extent to which firms will want to respond to the information about firm-specific shocks contained in the previous period's signals, corresponding to our previous parameter γ. The rationale for the parametrization of idiosyncratic shocks in terms of transitory and persistent components will be discussed further, when we turn to our quantitative results.

[7] Therefore, if we allowed firms to observe the current signals in real time while making decisions (as would be the case in a noisy rational expectations equilibrium), the demand and wage observations would be sufficient to exactly infer the current full information targets, and as a consequence the equilibrium would lead to full price adjustment.

Finally, we assume that, at time t, the shocks $(u_{t-T}, v^1_{i,t-T}, \tilde{v}^1_{t-T}, v^2_{i,t-T}, \tilde{v}^2_{t-T})$ become common knowledge (where T is large, but finite). As we will see in the next section, this assumption will help us cast the problem as a finite-dimensional filtering problem, and thereby lead to simple and tractable characterization of equilibrium dynamics.

In summary, a firm's information set at the beginning period t is defined as

$$\mathcal{I}^i_t = \left\{x_{it-s}, \omega_{it-s}, u_{t-T-s}, v^1_{i,t-T-s}, \tilde{v}^1_{t-T-s}, v^2_{i,t-T-s}, \tilde{v}^2_{t-T-s}\right\}_{s=1}^{\infty}.$$

Because states are fully revealed with a delay of T periods, the information contained in the demand and wage signals that are more than T periods old is redundant.

3.2 Solving the model

We solve this model using techniques and characterization results first suggested by Hellwig (2002) and further developed in Hellwig (2008a).[8] The central technical issue in this class of models is the infinite-regress problem that results from (a) the firm's desire to forecast the other firms' prices due to the complementarity in price-setting, and (b) their need to form a forecast of other firms' prices in order to interpret the demand signal x_{it}, when $\theta \neq \psi^{-1}$. In general, such higher-order filtering problems can quickly become intractable, since the model typically does not give rise to a recursive structure with a finite-dimensional state vector - in other words, Kalman filtering techniques, which in this class of models work well in single-agent filtering problems, are a lot less tractable for the higher-order forecasting problem that we face here. This problem is compounded by the endogeneity of information to equilibrium strategies through the demand signals.[9]

By assuming that fundamental shocks become common knowledge after a delay of T periods, we can recast the firms' decision problem as a filtering problem over the finite-dimensional vector of shocks, and thereby circumvent the tractability issues associated with the infinite regress problem. The resulting filtering problem is very tractable and admits a simple closed form solution, for a given information structure. The endogeneity of the demand signals to equilibrium prices is then resolved by finding a fixed point between the pricing conjecture entering the firm's filtering problem, and the resulting best-response prices.

[8] See also Hellwig and Veldkamp (2008) for similar equilibrium characterization results in a static model context.

[9] Woodford (2001) does solve his model using Kalman filtering techniques, but in his case, the information structure has no endogenous elements and admits a simple recursive representation of higher-order expectations.

Vector Representation and Model Solution: Let U_t, V_{it}^1, $\tilde{V}_{it}^1$ and V_{it}^2 and $\tilde{V}_{it}^2$ denote the vector of shocks that have occurred, but not yet been fully revealed at the time of choosing p_{it}; that is, $U_t' = (u_{t-1}, u_{t-2}, ...u_{t-T})$, $V_{it}^{1\prime} = (v_{it-1}^1, v_{it-2}^1, ...v_{it-T}^1)$, $\tilde{V}_{it}^{1\prime} = (\tilde{v}_{it-1}^1, \tilde{v}_{it-2}^1, ...\tilde{v}_{it-T}^1)$, $V_{it}^{2\prime} = (v_{it-1}^2, v_{it-2}^2, ...v_{it-T}^2)$, and $\tilde{V}_{it}^{2\prime} = (\tilde{v}_{it-1}^2, \tilde{v}_{it-2}^2, ...\tilde{v}_{it-T}^2)$.[10]

We separate the firms' optimal target prices and signals into two components - a 'common knowledge component' which consists of the contributions of all the shocks realized prior to date $t-T$, which are therefore common knowledge at date t, and a 'filtering component', which consists of the contributions of all the more recent shocks, i.e. the vectors U_t, V_t^1 and V_t^2. Formally, from the firms best response function (11), and the MA representation of the shocks (12), we define the common knowledge component of the firm's optimal price as

$$\hat{p}_{it} = \frac{Const}{1+\theta\kappa} + \frac{1}{1+\theta\kappa} \sum_{s=T+1}^{\infty} \left(\kappa\sigma_b \rho_b^s v_{i,t-s}^1 + \sigma_z \rho_z^s v_{i,t-s}^2 \right) + m_{t-T-1},$$

where we have already made use of the fact that the common knowledge component of the aggregate price, $\hat{p}_t$, equals m_{t-T-1}. The firms' optimal price is the sum of the commonly known component and the part that depends on signals from the last $T-1$ periods.

$$p_{it} = \hat{p}_{it} + (1-r)\,\sigma_u \mathbf{1}'\mathbb{E}_{it}[U_t] + r\left(\mathbb{E}_{it}[p_t] - \hat{p}_t\right) + \sigma_b \underbrace{\frac{\kappa\rho_b}{1+\theta\kappa}\boldsymbol{\Upsilon}_{\mathbf{b}}{}'}_{\gamma_{\mathbf{b}}'} \mathbb{E}_{it}[V_{it}^1] + \sigma_z \underbrace{\frac{\rho_z}{1+\theta\kappa}\boldsymbol{\Upsilon}_{\mathbf{z}}{}'}_{\gamma_{\mathbf{z}}'} \mathbb{E}_{it}[V_{it}^2] \quad (15)$$

where $\boldsymbol{\Upsilon}_{\mathbf{b}}{}' \equiv (1, \rho_b, \rho_b^2, ..., \rho_b^{T-1})$, $\boldsymbol{\Upsilon}_{\mathbf{z}}{}' \equiv (1, \rho_z, \rho_z^2, ..., \rho_z^{T-1})$ and $\mathbf{1} = (1, 1, ..., 1)'$, and $r = \frac{\kappa\left(\theta - \psi^{-1}\right)}{1+\theta\kappa}$. Note that this is in exactly the same form as the optimal pricing rule (2) from the simple example in Section 2. The vectors $\gamma_{\mathbf{b}}'$ and $\gamma_{\mathbf{z}}'$ are analagous to the parameter γ in the example.

Next, we conjecture that equilibrium prices will fully adjust to the shocks included in the common knowledge component, but the response to shocks included in the filtering component will be determined from the resulting equilibrium filtering problem. In other words, we conjecture that

$$p_t = \hat{p}_t + \sigma_u \phi' U_t \quad (16)$$

for some $T \times 1$ vector ϕ'=$(\phi_1, ...\phi_T)$. With this conjecture, we arrange firm i's optimal price as

$$p_{it} = \hat{p}_{it} + \sigma_u \left[(1-r)\,\mathbf{1}' + r\phi'\right] \mathbb{E}_{it}[U_t] + \sigma_b \gamma_{\mathbf{b}}' \mathbb{E}_{it}[V_{it}^1] + \sigma_z \gamma_{\mathbf{z}}' \mathbb{E}_{it}[V_{it}^2] \quad (17)$$

[10]Notice that we do not include u_t, v_{it}^1, $\tilde{v}_{it}^1$, v_{it}^2 and $\tilde{v}_{it}^2$ in these vectors as these shocks have not yet been realized. They are certainly relevant for the firms' optimal prices when choosing p_{it}, but since the firms have no information about these shocks yet, they have an expected value of zero and therefore no effect on the optimal price.

The firms' optimal pricing decisions thus depend on its expectations about the aggregate component U_t, as well as its expectations about the firm-specific shocks V_{it}^1 and V_{it}^2. Averaging (17) across all i and using the equilibrium conjecture $p_t = \hat{p}_t + \sigma_u \phi' U_t$, we find

$$\sigma_u \phi' U_t = \sigma_u \left[(1-r) \mathbf{1}' + r\phi' \right] \bar{E}_t[U_t] + \sigma_b \gamma_{\mathbf{b}}' \bar{E}_t[V_{it}^1] + \sigma_z \gamma_{\mathbf{z}}' \bar{E}_t[V_t^2], \tag{18}$$

where the $\bar{E}_t(\cdot) = \int E_{it}[\cdot] di$ denotes the firm's average expectations at the time of making the period t pricing decision. We thus obtain a representation for average prices as a function of average expectations about the underlying shocks. Average prices respond to U_t for two reasons: (i) the firms' average expectations about aggregate conditions (the first term), and (ii) the firms' average expectations of firm-specific shocks (the second and third terms). Our equilibrium conjecture is confirmed if (as will be shown below) these average expectations are linear functions of U_t.

We make use of the underlying information structure to compute these average expectations. Let X_{it} and Ω_{it} be the set of non-redundant signals at time t: $X_{it}' = (x_{i,t-1}, x_{i,t-2}, ...x_{t-T})$ and $\Omega_{it}' = (\omega_{i,t-1}, \omega_{i,t-2}, ...\omega_{t-T})$. Just like the target prices, we decompose the signal vectors X_{it} and Ω_{it} into a common knowledge component and a filtering component. For any vector $d' = (d_1, d_2, ...d_T)$, let $B(d)$ denote the upper-dimensional $T \times T$ matrix with d_i in its $k, k+i-1$-th entry:

$$B(d) \equiv \begin{pmatrix} d_1 & d_2 & . & d_{T-1} & d_T \\ 0 & d_1 & d_2 & . & d_{T-1} \\ . & . & . & . & . \\ 0 & . & 0 & d_1 & d_2 \\ 0 & 0 & . & 0 & d_1 \end{pmatrix}$$

With this notation, the signal vectors are written as

$$\begin{aligned} X_{it} &= \hat{X}_{it} + \sigma_u \left(\psi^{-1} B(\mathbf{1}) + \left(\theta - \psi^{-1}\right) B(\hat{\phi}) \right) U_t + \sigma_b B(\mathbf{\Upsilon_b}) V_{it}^1 + \tilde{\sigma}_b \tilde{V}_{it}^1 \\ \Omega_{it} &= \hat{\Omega}_{it} + \sigma_u B(\mathbf{1}) U_t + \sigma_z B(\mathbf{\Upsilon_z}) V_{it}^2 + \tilde{\sigma}_z \tilde{V}_{it}^2 \end{aligned}$$

where

$$\hat{\phi}' \equiv (0, \phi_1, ..., \phi_{T-1})$$

and the common knowledge components are $\hat{X}_{it}' = (\hat{x}_{it-1,t}, ..., \hat{x}_{it-T,t})$ and $\hat{\Omega}_{it}' = (\hat{\omega}_{it-1,t}, ..., \hat{\omega}_{it-T,t})$, with

$$\begin{aligned} \hat{x}_{it-k,t} &= \sigma_b \sum_{s=T+1}^{\infty} \rho_b^{s-k} v_{i,t-s}^1 + \theta m_{t-T-1} \\ \hat{\omega}_{it-k,t} &= \sigma_z \sum_{s=T+1}^{\infty} \rho_z^{s-k} v_{i,t-T-s}^2 + m_{t-T-1}. \end{aligned}$$

The next lemma computes the average expectations of $\left(U_t, V_{it}^1, V_{it}^2\right)$, using the solution to the firms' filtering problem for $\left(U_t, V_{it}^1, V_{it}^2\right)$ conditional on (X_{it}, Ω_{it}):

Lemma 1 *(Average expectations lemma): (i) $\bar{E}_t[U_t]$, $\bar{E}_t[V_{it}^1]$, and $\bar{E}_t[V_{it}^2]$ are given by*

$$\begin{aligned}
\bar{E}_t[U_t] &= \Gamma' \left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1} \Gamma U_t, \\
\bar{E}_t[V_{it}^1] &= \begin{bmatrix} \sigma_b B(\mathbf{\Upsilon_b})' & 0 \end{bmatrix} \left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1} \Gamma U_t, \\
\bar{E}_t[V_{it}^2] &= \begin{bmatrix} 0 & \sigma_z B(\mathbf{\Upsilon_z})' \end{bmatrix} \left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1} \Gamma U_t
\end{aligned}$$

where

$$\Gamma = \sigma_u \begin{pmatrix} \psi^{-1} B(\mathbf{1}) + \left(\theta - \psi^{-1}\right) B(\hat{\phi}) \\ B(\mathbf{1}) \end{pmatrix} \text{ and } \Delta = \begin{pmatrix} \sigma_b B(\mathbf{\Upsilon_b}) & \tilde{\sigma}_b I & 0 & 0 \\ 0 & 0 & \sigma_z B(\mathbf{\Upsilon_z}) & \tilde{\sigma}_z I \end{pmatrix}.$$

To interpret these expressions, notice that the firms' posterior variance-covariance matrix over U_t, conditional on (X_{it}, Ω_{it}), is

$$\Sigma = I - \Gamma' \left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1} \Gamma.$$

The average expectations $\bar{E}_t[U_t]$ can thus be rewritten as $\bar{E}_t[U_t] = [I - \Sigma] U_t$. Similarly, the two terms $\begin{bmatrix} \sigma_b B(\mathbf{\Upsilon_b})' & 0 \end{bmatrix} (\Gamma\Gamma' + \Delta\Delta')^{-1} \Gamma$ and $\begin{bmatrix} 0 & \sigma_z B(\mathbf{\Upsilon_z})' \end{bmatrix} (\Gamma\Gamma' + \Delta\Delta')^{-1} \Gamma$ in turn denote the negative of the conditional covariance matrices of V_{it}^1 and V_{it}^2 with U_t, respectively.

The proof of this lemma is straight-forward: using standard projection arguments, we first derive expectations for U_t, V_{it}^1 and V_{it}^2 conditional on (X_{it}, Ω_{it}). These are then averaged across all individuals to find average expectations of U_t in terms of the signal averages - the signal averages in turn are just functions of the vector of aggregate shocks U_t. The matrix Γ and variance-covariance matrix Σ depend on $\hat{\phi}$, because the demand signals depend on aggregate price adjustment vector $\hat{\phi}$.

Substituting this characterization of average expectations into (18), we obtain, for given equilibrium conjecture $\hat{\phi}$ and the resulting posterior variance-covariance matrix Σ, the vector ϕ that result from (18). An equilibrium is then a fixed point of (18), given the characterization of average expectations in Lemma 1. For given parameters, the solution to this fixed point problem is easy to compute numerically, even for large values of T.

Analytical Results: Before presenting numerical results, we can provide analytical results for some special cases. When $\rho_b = \rho_z = 0$, only the expectations about aggregate shocks matter for price adjustment - firm-specific shocks do not affect optimal prices. This is the case on which existing results about propagation and persistence are based. In this case, ϕ' is given by

$$\phi' = (1-r)\,\mathbf{1}'\,(I-\Sigma)\,[I-r\,(I-\Sigma)]^{-1}\,. \tag{19}$$

When $r=0$, the firms' prices then adjust according to their expectations of U_t, i.e. $p_t = \mathbf{1}'\bar{E}_t[U_t]$. When $r>0$, since $(1-r)\,[I-r\,(I-\Sigma)]^{-1} = (1-r)\sum_{s=0}^{\infty} r^s\,(I-\Sigma)^s << I$, real effects will be larger. Therefore, we have the standard effect that complementarities delay price adjustment. This is the generalization of proposition 1 from the example section, and corresponds to the main result of Woodford (2001).[11]

When instead the firm-specific shocks are persistent ($\rho_b > 0$ and/or $\rho_z > 0$), the firms' optimal prices respond to their expectations about idiosyncratic conditions, i.e. the second and third terms in (18) are positive. This will increase the firms' overall price adjustment, which is now characterized as

$$\begin{aligned}\phi' &= (1-r)\,\mathbf{1}'\,(I-\Sigma)\,[I-r\,(I-\Sigma)]^{-1} \\ &\quad + \frac{1}{\sigma_u}\left[\begin{array}{cc}\sigma_b^2\gamma_{\mathbf{b}}' B(\mathbf{\Upsilon_b})' & \sigma_z^2\gamma_{\mathbf{z}}' B(\mathbf{\Upsilon_z})'\end{array}\right]\left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1}\Gamma\,[I-r\,(I-\Sigma)]^{-1}\,.\end{aligned} \tag{20}$$

The relative sizes of aggregate, as well as persistent and transitory idiosyncratic shocks determines the magnitudes of the two components to overall price adjustment. When aggregate shocks are large relative to idiosyncratic ones, the posterior variance-covariance matrix Σ will be small, and adjustment dynamics will mainly be driven by the first term which captures the firms' expectations about aggregate shocks - firms adjust prices for the 'right' reasons.

When instead aggregate shocks are small relative to idiosyncratic ones, Σ is close to I, firms do not update much their beliefs about aggregates and instead attribute all changes in signals to idiosyncratic shocks. As a consequence, the first term in (20) is small. The response of prices to aggregate shocks is then driven mainly by the second term in (20), i.e. prices adjust 'for the wrong reasons', because firms mistakenly attribute changes in aggregate conditions to changes in idiosyncratic shocks. The size of this second term depends on the conditional co-variances between aggregate and persistent idiosyncratic shocks. The size of these covariances in turn depends on the relative importance of transitory and persistent idiosyncratic shocks.

We can make these observations precise by considering the limiting case where $\sigma_u \to 0$, hold-

[11] In addition, the endogeneity of demand signals generates extra persistence, because monetary shocks are only gradually reflected in demand signals.

ing fixed the other parameters. In that case, $\Sigma \to I$, and ϕ converges to

$$\begin{aligned}\phi' &= \mathbf{e}_1'\rho_b\sigma_b^2 B(\boldsymbol{\Upsilon}_\mathbf{b})B(\boldsymbol{\Upsilon}_\mathbf{b})'\left[\tilde{\sigma}_b^2 I + \sigma_b^2 B(\boldsymbol{\Upsilon}_\mathbf{b})B(\boldsymbol{\Upsilon}_\mathbf{b})'\right]^{-1}\frac{\kappa}{1+\theta\kappa}\left[\psi^{-1}B(\mathbf{1}) + \left(\theta - \psi^{-1}\right)B(\hat{\phi})\right] \\ &\quad + \mathbf{e}_1'\rho_z\sigma_z^2 B(\boldsymbol{\Upsilon}_\mathbf{z})B(\boldsymbol{\Upsilon}_\mathbf{z})'\left[\tilde{\sigma}_z^2 I + \sigma_z^2 B(\boldsymbol{\Upsilon}_\mathbf{z})B(\boldsymbol{\Upsilon}_\mathbf{z})'\right]^{-1}\frac{1}{1+\theta\kappa}B(\mathbf{1}),\end{aligned}$$

where $\mathbf{e}_1' = (1, 0, ..., 0)$. In this limiting case, the firms' price response depends on (i) the serial correlation of the persistent idiosyncratic shocks, as determined by ρ_b and ρ_z, and (ii) relative importance of the transitory vs. persistent idiosyncratic shocks to demand and costs, as determined by $\sigma_b^2 B(\boldsymbol{\Upsilon}_\mathbf{b})B(\boldsymbol{\Upsilon}_\mathbf{b})'$ vs. $\tilde{\sigma}_b^2 I$ and $\sigma_z^2 B(\boldsymbol{\Upsilon}_\mathbf{z})B(\boldsymbol{\Upsilon}_\mathbf{z})'$ vs. $\tilde{\sigma}_z^2 I$. Our general model thus replicates the findings of the example and confirms our earlier interpretation for the example's reduced form parameters.

As with the example, the resulting price adjustment 'for the wrong reasons' can be substantial, and in some cases, converge to complete adjustment. This occurs, in particular, if there are no transitory idiosyncratic shocks so that firms attribute all fluctuations in signals to the persistent firm-specific cost and demand conditions. The following proposition characterizes the adjustment dynamics if in addition to $\sigma_u \to 0$, we set $\tilde{\sigma}_b = \tilde{\sigma}_z = 0$:

Proposition 3 *(Setting the right prices for the wrong reasons): In the limit as $\sigma_u \to 0$, and $\tilde{\sigma}_b = \tilde{\sigma}_z = 0$, the impulse response function to nominal shocks is given by the vector ϕ, where*

$$\phi_s = \left(\frac{\kappa\psi^{-1}\rho_b + \rho_z}{1+\theta\kappa}\right)\frac{1-(r\rho_b)^s}{1-r\rho_b} \tag{21}$$

denote the impact of u_{t-s} on p_t.

Thus, average prices adjust even though firms remain completely in the dark about the aggregate shocks. In this particular case, the adjustment takes the form of a geometric convergence at rate $r\rho_b$ to a permanent adjustment level given by

$$\lim_{s\to\infty}\phi_s = \left(\frac{\kappa\psi^{-1}\rho_b + \rho_z}{1+\theta\kappa}\right)\frac{1}{1-r\rho_b}.$$

The persistence in the adjustment in this particular case results from the fact that aggregate shocks affect demand observations only gradually, and with delay: Demand will respond by $(1-r)$ on impact, but prices do not change. The following period, prices respond to the original increase, which further raises the demand signal, and so on. Over time, aggregate prices then gradually increase to the permanent adjustment level, as the shock gets reflected in aggregate prices, and hence in demand. When $\rho_b = \rho_z = \rho$, we can rearrange this as

$$\phi_s = \frac{\rho(1-r)}{1-r\rho}\left(1-(r\rho)^s\right)$$

and we therefore observe that the permanent adjustment level $\rho\,(1-r)\,/\,(1-r\rho)$ is an increasing function of the degree of persistence, and a decreasing function of the complementarity r. When $r = 0$, demand observations are not affected by the aggregate price level, and the permanent adjustment level is reached after the first period.[12]

Moreover, when $\rho = 1$, prices eventually reach full adjustment, i.e. $\lim_{s\to\infty}\phi_s = 1$. In this particular case, firms will fully adjust to the idiosyncratic shocks over time. When an aggregate shock occurs, they will confuse it with an idiosyncratic one, but they nevertheless fully adjust their prices to these permanent shocks to their demand and costs.

4 Quantitative Results

In this section, we proceed to a quantitative evaluation of our model. In particular, we aim to explore the quantitative effects of information and payoff heterogeneity in our price-setting model.

The standard deviation of (monthly) innovations to money growth σ_u is set to $.0036$, following Golosov and Lucas (2007). The parameters θ, ψ and κ are borrowed from the literature and set to $\psi = 2, \theta = 4$ and $\kappa = 1$. These parameters imply that the pricing complementarity $r \equiv \kappa(\theta - \psi^{-1})/(1+\theta\kappa)$ is equal to 0.7. Thus, our calibration has a high degree of complementarity, which would suggest a slow response of prices to aggregate monetary shocks. However, as we will see, the effect of complementarities is muted by the presence of idiosyncratic shocks in the model.

The remaining parameters of the model relate to the stochastic process for the idiosyncratic shocks. Before explaining our calibration strategy, it will be useful to discuss how these parameters interact with each other. For this purpose, we shall, for the moment, set the variance of the two transitory idiosyncratic shocks to zero, and just focus on the case in which idiosyncratic demand and wages are AR(1) processes.

Figure 3 illustrates this interaction between the persistence of idiosyncratic shocks and the degree of complementarity. In the panel on the right, ρ_b and ρ_z are close to 1, so the impulse response functions show a rapid adjustment to monetary shocks, almost independently of the degree of complementarity. When the persistence is lower, as in the panel on the left, complementarity plays a bigger role in determining the speed of adjustment. In both figures, prices rapidly adjust on impact from the firm-specific component, followed by slow adjustment as the firms

[12]The issue here is not the complementarity per se, but the fact that the same parameters that determine r also determine the effects of aggregate prices on the demand signal. Since prices respond to demand signal with a one period delay, demand signals also incorporate the aggregate shock only gradually, as it enters into prices.

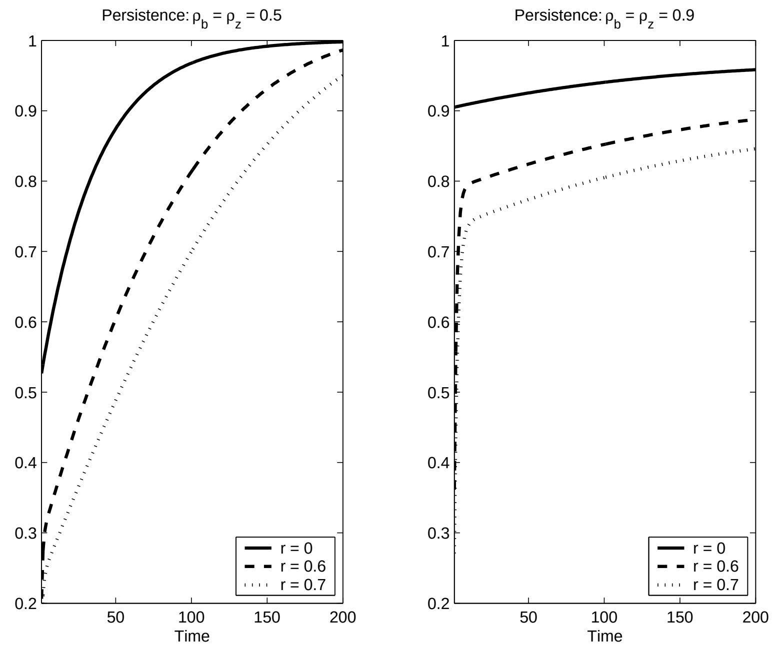

Figure 3: Impulse Responses to a monetary shock

fully separate idiosyncratic from aggregate shocks.[13] The complementarity affects both the initial adjustment level and the subsequent speed of convergence, but its effects are stronger when the firm-specific shocks are less persistent.

Figure 4 highlights a similar interaction between the persistence parameters and the effect of informational heterogeneity[14]. When ρ_b and ρ_z are close to 1, prices respond very quickly to monetary shocks. This is true even if informational heterogeneity is large i.e. even if the variance of the idiosyncratic shocks is high relative to the monetary shock. This is evident from the panel on the right, which plots the impulse response functions for various values of σ_b^2 and σ_z^2. In the other panel, shocks are much less persistent and changes in informational heterogeneity have much larger effects.

Finally, Table 1 shows the implications of the persistence and variance parameters for price

[13]Higher persistence leads to higher responsiveness in the short to medium term, but it also makes the filtering problem more challenging in the long run. So, the convergence to full adjustment takes much longer with persistent idiosyncratic shocks. This conclusion is not robust to adding additional sources of information over the long run. Recall that firms in this model have access to a very limited set of signals. This restriction seems reasonable over the short-to-medium term, but less so if we are interested in long-term convergence. Since our focus is on the former horizon, we do not address this feature of the learning process in this paper.

[14]In order to isolate the effects of heterogeneity, we set the degree of complementarity to zero.

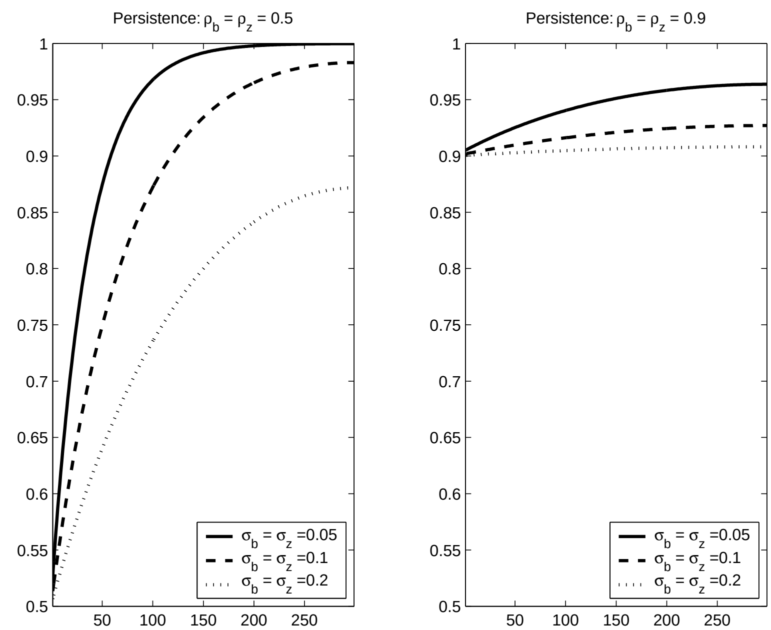

Figure 4: Impulse Responses to a monetary shock

dispersion. In particular, note that for a given level of the variances σ_b^2 and σ_z^2, higher persistence increases price dispersion. However, as we saw earlier, higher persistence also implies quicker adjustment to monetary shocks. Thus, the only way to have both reasonable price dispersion and persistent real effects of monetary shocks is to have high variances for the idiosyncratic shock processes.

	$\sigma_b = \sigma_z = 0.02$	$\sigma_b = \sigma_z = 0.05$	$\sigma_b = \sigma_z = 0.10$
$\rho_b = \rho_z = .5$	0.003	0.009	0.017
$\rho_b = \rho_z = .7$	0.005	0.014	0.028
$\rho_b = \rho_z = .9$	0.012	0.029	0.058

Table 1: Price dispersion for different levels of persistence and std. deviations of idiosyncratic shocks

4.1 Calibrating the remaining parameters

The earlier discussion shows that the quantitative implications of the model for the effect of monetary shocks are affected significantly by the parameters governing the stochastic processes for the idiosyncratic shocks. To discipline our choice of these parameters, we target four key moments

1. Price dispersion (the standard deviation of the log of relative prices) of 6-10%

2. Quantity dispersion (the standard deviation of log of quantities) of 25-30 %

3. Correlation between prices and quantities of roughly -0.20

4. Autocorrelation (daily) of the log of relative prices of 0.98.

The dispersion and autocorrelation of relative prices are related to the persistence and variance of the idiosyncratic shock processes. Our target for price dispersion is derived from the statistics reported by Burstein and Hellwig (2007) for the Dominicks scanner price data.[15]

The autocorrelation target is derived from the monthly number estimated by Midrigan (2007) for the same data. Recall that this parameter plays a crucial role in determining the role of the idiosyncratic pricing motive and thereby, has important implications for the responsiveness of prices. Two issues arise in picking a number for this target. First, several studies have found substantial evidence of nominal stickiness in the data. In our model, on the other hand, prices are fully flexible. Given this feature of our model, we are probably overstating the autocorrelation in shocks by matching the reported autocorrelation numbers too closely. However, it turns out that our quantitative results and main conclusions are insensitive to small changes in this autocorrelation target.

The second issue is related to the choice of period length. We would like to think of our model as a close proxy for an environment of real-time information availability. In light of this, our preferred period length is a day. Since Midrigan (2007) reports a monthly number for the autocorrelation of relative prices, we have to adjust it to get a daily number. One simple way to do this to assume an AR(1) process for relative prices. In this case, a monthly autocorrelation estimate of 0.65 implies a daily autocorrelation of 0.98. However, as the following discussion will

[15]Burstein and Hellwig find price dispersion measures of roughly 10%. We did not find direct measures of relative price dispersion in papers using other data sources, but this level seems consistent with the widely reported numbers on the magnitude of price changes (Klenow and Kryvtsov(2008), Bils and Klenow (2004), Nakamura and Steinsson (2008)).

show, adjusting period length in the model has implications for the amount of information flow. This is an important issue and we discuss it below in greater depth.

Finally, the quantity moments targeted in the calibration help us identify the demand shock parameters separately from the cost shock ones. To see how, note that, from the demand function of the firm

$$c_{it} - c_t = b_{it} - \theta(p_{it} - p_t) \tag{22}$$

Thus, both shocks affect relative quantities through their effect on relative prices, but quantities are also directly affected by the demand shock. In a model with only cost shocks, (22) implies that relative quantities and prices would would be perfectly negatively correlation. This is at odds with the data. Burstein and Hellwig (2007), for example, find that the correlation is only modestly negative in the Dominick's data. This feature of the data, they argue, suggests that both demand and cost shocks are necessary to simultaneously explain the observed price and quantity moments. Eichenbaum et al.(2008) support similar observations about demand fluctuations using a different, much wider data set.

4.2 Effect of Period Length

The choice of the length of a period has implications for 2 key aspects of the model. First, as mentioned earlier, it requires an adjustment to our targets for autocorrelation of prices. In general, less mean reversion from one model period to the next will generate more price adjustment. Second, it changes the sampling frequency and through that, the rate of arrival of new information. Shortening the period length increases the per-period flow of information and therefore, allows firms to learn faster.

To see these effects more clearly, consider the following 2 versions of the model. The first version has the length of a period set to a day, while the second one uses a month for period length. The variance of the monetary process and the target for the autocorrelation of prices are scaled accordingly. Again, we assume that there are no transitory shocks to idiosyncratic shock processes (i.e. they consist of only the AR(1) components). The parameters of the AR(1) are then set so as to match the target moments listed above. Table 2 shows us a few values for the impulse-response functions. As the table shows, the daily model shows a much quicker adjustment of prices to monetary shocks. This difference persists even when we look at responses a year after the shock. What explains this stark difference in the two models? First, firms in the daily model draw signals more frequently (30 times a month, as opposed to once a month in the monthly

	Impact	1 mth	6 mths
Daily Model	0.29	0.91	0.91
Monthly Model	0.19	0.37	0.40

Table 2: Impulse Responses to a monetary shock

model) and are therefore able to learn about the shocks much faster. Moreover, the smaller degree of mean reversion from one period to the next also makes optimal prices more responsive. Both of these effects lead to a quicker response of prices.

This simple exercise suggests that if we are allowed to set the period length without controlling the information flow, the model can generate any degree of price adjustment, ranging from almost perfectly responsive (very short period lengths) to significant rigidity (with very long periods).

We can partly address this problem by using the transitory components of the idiosyncratic shock process. In particular, by increasing the variance of these noise terms in a model with shorter period lengths, we can slow down the learning process. We follow this approach in our benchmark calibration for the daily model and set the standard deviations of this transitory component so that the speed of learning (as measured by the variance of the posterior distribution for the persistent component) is the same as in the monthly model.

4.3 Results

The impulse response functions of our benchmark daily model are shown in Figure 5. The addition of transitory noise to the signals increases the rigidity on impact (the aggregate price level reflects only 20 % of the aggregate shock on impact), but this effect dies out very quickly. In little over a month, prices have incorporated 90% of the shock. These results show that the tension highlighted in the simple example at the beginning of the paper - between generating both meaningful delays in price adjustment and realistic levels of price dispersion - survives even in a micro-founded model.

That said, the graph also reveals that complete convergence takes a long time. In other words, the real effects of a monetary shock, albeit small in magnitude, have a highly persistent component. When both idiosyncratic and aggregate shocks are highly persistent, it takes the firms a very large number of periods to separate the two using the market-generated signals. However, since the desired responses to an aggregate vs. an idiosyncratic shock are only slightly different, the

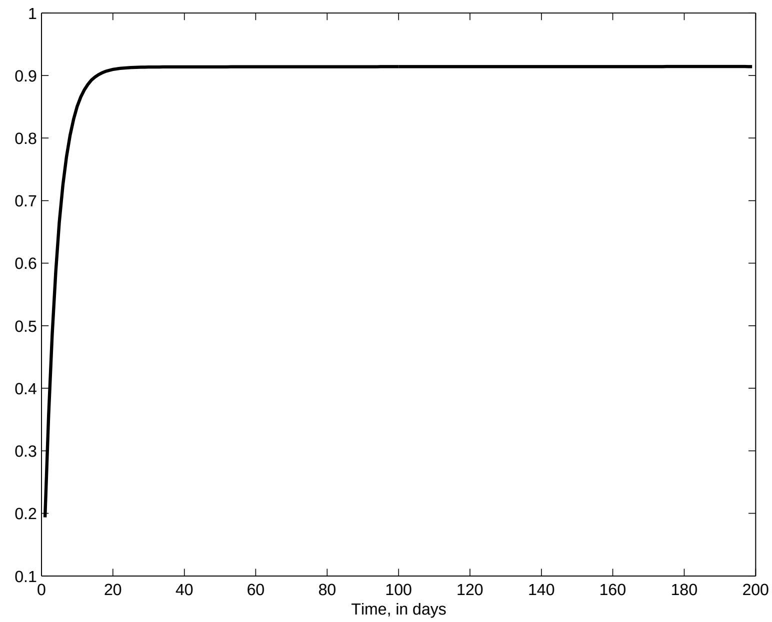

Figure 5: Impulse responses to a monetary shock in the benchmark daily model

firms' optimal response is to make a fairly rapid adjustment to almost the right level and then wait for further information before full convergence.

5 Extensions

In this section, we test the robustness of our conclusions through a couple of extensions to our benchmark analysis from the previous section. The first one adds productivity shocks - both aggregate and idiosyncratic - to the model, while the second explores the effect of larger aggregate shocks on our results. We show that both of these versions face the same (or in some cases, an even larger) difficulty in generating persistent real effects for monetary shocks, while matching micro moments.

5.1 Productivity Shocks

The purpose of this extension is to look at the effects of adding another source of price dispersion. We show numerically that by separating the processes of learning about cost and demand-side shocks, we can generate both reasonable levels of price dispersion and slow adjustment to monetary shocks. We illustrate this point by replicating the results of Mackowiak and Wiederholt (2008), who study a model of price adjustment with rational inattention, in which firms have to divide a limited information processing capacity between tending to idiosyncratic technology and aggregate demand shocks.[16] Since idiosyncratic technology shocks are an order of magnitude larger than aggregate demand shocks, firms tend to pay very little attention to the latter, so that prices appear to be sticky in response to aggregate demand shocks, yet very responsive to firm-specific technology shocks. However, we will see that these results have the strongly counterfactual prediction that prices and quantities are perfectly negatively correlated.

We incorporate productivity shocks by assuming that the log of the productivity term in (9) is the sum of an aggregate and an idiosyncratic component:

$$a_{it} = g_t + g_{it} \tag{23}$$

where g_t is an aggregate term and g_{it} is a firm-specific idiosyncratic factor. Both these components in turn are modeled as AR(1) processes:

$$g_t = \sigma_A \sum_{s=0}^{\infty} \rho_A^s u_{t-s}^2 \qquad \text{and} \qquad g_{it} = \sigma_a \sum_{s=0}^{\infty} \rho_a^s v_{i,t-s}^3 \tag{24}$$

where $u_t^2, v_{i,t}^3$ are distributed according to $\mathcal{N}(0,1)$ and $\chi, \xi, \sigma_a, \sigma_A$ are non-negative.[17]

Firms observe a_{it} but not the aggregate and idiosyncratic components separately. As before, firms make their choices in period t before observing the shocks for that period. The other elements of the information structure are the same as in the benchmark model. In particular, productivity shocks become common knowledge T periods after they occur.

The FOC of the firm's price-setting problem takes the form

$$P_{it}^{1+\theta\kappa} = \frac{\theta(1+\kappa)\,\hat{\imath}}{\theta-1} \cdot \frac{\mathbb{E}_{it}[Z_{it} B_{it}^{1+\kappa} A_{it}^{-1-\kappa} C_t^{1+\kappa} P_t^{\theta(1+\kappa)}]}{\mathbb{E}_{it}[B_{it} C_t P_t^{\theta} M_t^{-1}]}$$

Log-linearizing the FOC yields

$$p_{it} = Const + \frac{\kappa}{1+\theta\kappa}\mathbb{E}_{it}[b_{it} + c_t + \theta p_t] + \frac{1}{1+\theta\kappa}\mathbb{E}_{it}[z_{it} + m_t] - \frac{1+\kappa}{1+\theta\kappa}\mathbb{E}_{it}[a_{it}] \tag{25}$$

Now, using the household FOC (5) to substitute for c_t, we have

$$p_{it} = Const + \frac{1}{1+\theta\kappa}\mathbb{E}_{it}[\kappa b_{it} + z_{it}] + \frac{1+\kappa\psi^{-1}}{1+\theta\kappa}\mathbb{E}_{it}[m_t] + \frac{\kappa\left(\theta - \psi^{-1}\right)}{1+\theta\kappa}\mathbb{E}_{it}[p_t] - \frac{1+\kappa}{1+\theta\kappa}\mathbb{E}_{it}[a_{it}] \tag{26}$$

[16] Our transitory demand and cost shocks can also be reinterpreted in terms of channel noise in a rational inattention model. In that case, the variances of the transitory shocks would be the result of a firms' optimal decision, subject to a capacity constraint on information processing.

[17] It is straight-forward, but not necessary for our purposes, to include transitory as well as persistent idiosyncratic shocks to technology.

As before, denote by U_t^2 and V_{it}^3 the vector of 'relevant' aggregate and idiosyncratic productivity shocks i.e. those that have been realized but have not yet been fully revealed at the time of making period t decisions. Then, as before, we write the optimal price in terms of a common knowledge component and the firms' conditional expectations of the shocks that havent yet been revealed:

$$\begin{aligned} p_{it} &= \hat{p_{it}} + \sigma_u (1-r)\, \mathbf{1}'\mathbb{E}_{it}[U_t] + r\left(\mathbb{E}_{it}[p_t] - \hat{p}_t\right) \\ &+ \sigma_b \gamma_b' \mathbb{E}_{it}[V_{it}^1] + \sigma_z \gamma_z' \mathbb{E}_{it}[V_{it}^2] \\ &- \sigma_A \underbrace{\rho_A \left(\frac{1+\kappa}{1+\theta\kappa}\right) \mathbf{\Upsilon_A'}}_{\gamma_{\mathbf{A}}'} \mathbb{E}_{it}[U_t^2] - \sigma_a \underbrace{\rho_a \left(\frac{1+\kappa}{1+\theta\kappa}\right) \mathbf{\Upsilon_a'}}_{\gamma_{\mathbf{a}}'} \mathbb{E}_{it}[V_{it}^3] \end{aligned} \tag{27}$$

where $\mathbf{\Upsilon_A'} = (1, \rho_A, \rho_A^2, ..., \rho_A^{T-1})$, $\mathbf{\Upsilon_a'} = (\mathbf{1}, \rho_{\mathbf{a}}, \rho_{\mathbf{a}}^{\mathbf{2}}, ..., \rho_{\mathbf{a}}^{\mathbf{T-1}})$.

The solution proceeds as in the benchmark model. We start with a conjecture about the aggregate price index

$$p_t = \hat{p}_t + \sigma_u \phi_1' U_t + \sigma_A \pi' U_t^2 \tag{28}$$

where $\phi' = (\phi_1, \phi_2, ...\phi_T)$ and $\pi' = (\pi_1, \pi_2, ...\pi_T)$ are real-valued vectors. Given this conjecture, the signal vectors are related to the realizations of the shocks as follows:

$$\begin{aligned} X_{it} &= \hat{X}_{it} + \sigma_u \left(\psi^{-1} B(\mathbf{1}) + \left(\theta - \psi^{-1}\right) B(\hat{\phi})\right) U_t + \sigma_A \left(\theta - \psi^{-1}\right) B(\hat{\pi}) U_t^2 + \sigma_b B(\mathbf{\Upsilon_b}) V_{it}^1 \\ \Omega_{it} &= \hat{\Omega}_{it} + \sigma_u B(\mathbf{1}) U_t + \sigma_z B(\mathbf{\Upsilon_z}) V_{it}^2 \\ A_{it} &= \hat{A}_{it} + \sigma_A B(\mathbf{\Upsilon_A}) U_t^2 + \sigma_a B(\mathbf{\Upsilon_a}) V_{it}^3, \end{aligned}$$

where A_{it} denotes the vector of productivity observations and $\hat{A}_{it}$ its common knowledge component. Note that the aggregate productivity shocks enter the demand signals through their effect on the aggregate price. Thus, the variance-covariance matrices depend jointly on ϕ and π. Finally, using the conjecture and averaging the FOC across firms, we get

$$\begin{aligned} \sigma_u \phi' U_t + \sigma_A \pi' U_t^2 &= \sigma_u \left[(1-r)\, \mathbf{1}' + r\hat{\phi}'\right] \bar{E}_t[U_t] - \sigma_A \left[\gamma_A' - r\hat{\pi}'\right] \bar{E}_t[U_t^2] \\ &+ \sigma_b \gamma_{\mathbf{b}}' \bar{E}_t[V_{it}^1] + \sigma_z \gamma_{\mathbf{z}}' \bar{E}_t[V_t^2] - \sigma_a \gamma_a' \bar{E}_t[V_{it}^3] \end{aligned} \tag{29}$$

As before, we can numerically solve the fixed point problem implicit in the above expression after characterizing the relevant average expectation terms from the firms' filtering problem.

We replicate the conclusions in Mackowiak and Wiederholt (2008) by (i) making the firm-specific wage and demand shocks completely transitory, so that they have no effect on firms' optimal prices and serve purely as informational noise, and (ii) using idiosyncratic technology shocks to match the observed price dispersion and persistence.

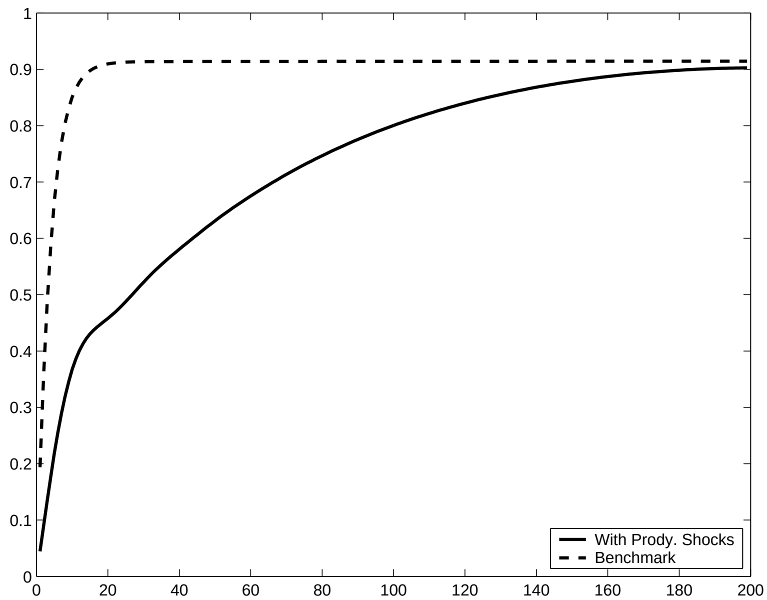

Figure 6: Impulse responses to a monetary shock with productivity shocks and transitory wage/demand shocks

Figure 6 plots the impulse response functions. As the graph shows, aggregate prices respond much more slowly to monetary shocks in this version compared to our benchmark model. As before, idiosyncratic cost and demand shocks are much larger than aggregate monetary shocks so firms tend to attribute the shocks largely to idiosyncratic factors. However, in contrast to our benchmark model, firms do not want to adjust their prices in response to these transitory shocks. In other words, firms choose to wait for more draws to learn more precisely about the aggregate shock. Since pricing complementarity is relatively high, firms also find it optimal to delay their responses to the perceived aggregate shock until other firms respond. As a result, aggregate prices respond rather sluggishly to small aggregate shocks. Yet at the same time, prices adjust rapidly to firm-specific technology shocks, which account for the large degree of price dispersion. This conclusion is similar to the findings in Mackowiak and Wiederholt (2008).

The row labeled Model I in Table 3 shows the relevant moments for this version of our model. Given our calibration, the model has the counter-factual implication that prices and quantities are highly negatively correlated. To see why this is so, recall from (22) that relative quantities are a linear combination of the demand shocks and a term involving relative price dispersion. In our calibration, demand shocks are assumed to completely transitory, so their contribution to overall dispersion of quantities is very small, so the dispersion in relative quantities comes almost exclu-

	Prc Disp	Prc Autocor	Prc-Qty Corr	Qty Disp
Targets	0.06-0.10	0.98	-0.20	0.25-0.30
Model I	0.08	0.98	-0.99	0.32
Model II	0.08	0.98	-0.35	1.05

Table 3: Relative Price and Quantity Moments with Productivity Shocks

sively from the dispersion in prices, leading to an almost perfect negative correlation. Of course, this suggests that we could reduce this correlation by increasing the variance of the demand shock process. The row labeled Model II in Table 3 is an example of such a calibration. While it fixes the correlation problem, it does so at the cost of counterfactually high quantity dispersion, which is at odds with the data. Thus, as in many standard sticky price models, demand shocks must play an important role if we are to simultaneously match quantity and price moments. In our environment, that means increasing the persistence of the shocks. But, this takes us closer to our benchmark model, where price adjustment occurs relatively fast, because firms wish to adjust prices in response to idiosyncratic demand shocks. In conclusion, this exercise illustrates that the basic tension faced by the benchmark model is not significantly mitigated even if we add productivity shocks. Moreover, it shows that the non-neutrality results in Mackowiak and Wiederholt rely on the separation of learning about idiosyncratic productivity and aggregate demand conditions.

5.2 Larger Aggregate Shocks

Next, we return to our benchmark model and consider the effect of larger aggregate shocks. One interpretation of this larger shock is to think of the firms in the model as belonging to the same sector of the economy and being subject to a sector-level demand shock in addition to the purely idiosyncratic ones. Many authors have found that sectoral demand shocks play an important role in explaining patterns in sectoral inflation. For example, De Gregorio et al. (1994) identify a demand shift towards nontradeables in the OECD data from 1970-85.

While a full-fledged multisector model is beyond the scope of this paper, we can gain some intuition for the effects of larger aggregate shocks on the learning process by increasing the variance of the aggregate shocks in our benchmark model. Figure 7 plots the impulse responses for various values of σ_u. To see the differences more clearly, we slow down the adjustment process by reducing persistence in the idiosyncratic shocks[18]. As the graph shows, higher values of the aggregate shock speeds up the adjustment process. The intuition for this comes from the fact as the variance of the aggregate shock increases, firms learn faster about the aggregate shock (because they are less likely to mistake it for an idiosyncratic one). This faster learning mutes the effect of the complementarity and the aggregate price responds quickly - firms set the right prices for the right reasons.

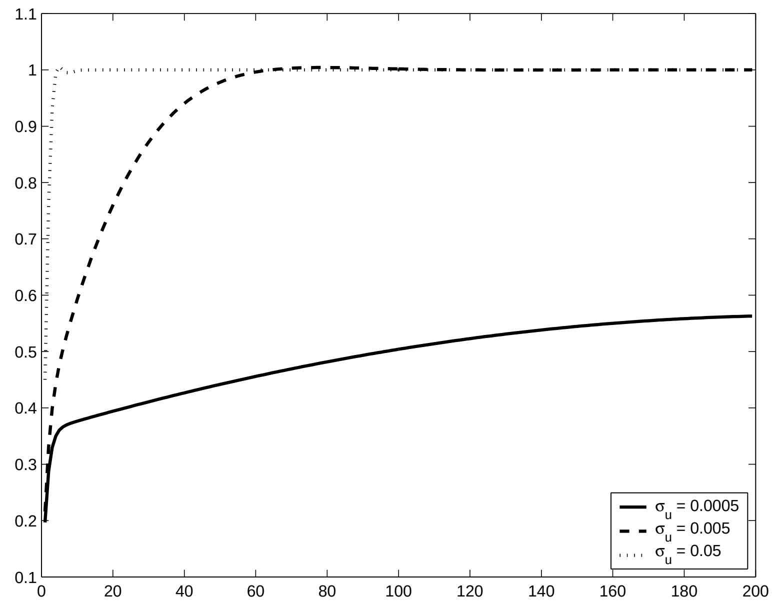

Figure 7: Impulse responses to a monetary shock for various levels of std deviation of aggregate shocks

6 Conclusion

In this paper, we have studied a model of price adjustment with heterogeneously informed firms, who face idiosyncratic as well as aggregate shocks, which they seek to infer from the information generated by their market activities. Two important conclusions emerge from our analysis.

First, the presence of transient firm-specific cost and demand shocks significantly mutes the real effects of monetary shocks on impact. Although these idiosyncratic shocks must be large to account for observed levels of price dispersion, thus generating large uncertainty about aggregate conditions, they also render the signal extraction problem less relevant for the firms' pricing de-

[18]This means that we will not be able to match our price dispersion and autocorrelation targets, but that does not affect our conclusions about the effects of larger shocks on learning.

cisions. As long as positive adjustment on impact is optimal in response to either shock, the firm will adjust its price, despite its confusion about the exact nature of the shock, i.e. idiosyncratic vs. aggregate. In the case where the shocks are equally persistent, prices quickly adjust to aggregate shocks despite the fact that firms remain permanently confused about the idiosyncratic and the aggregate fluctuations they are exposed to. This result of almost complete neutrality is reminiscent of related flexibility results in the sticky price literature (e.g. Caplin and Spulber (1987), or Golosov and Lucas (2007)), although the underlying reasons are very different.

Second, at longer horizons, our minimalist information structure, which is meant to capture the large degree of product-level dispersion, leads to implausible large delays in learning. While the response on impact is sizeable and positive, it is not complete, and in the absence of additional information, it takes firms a long time to completely adjust prices in response to permanent nominal shocks. We view this result as an artifact of our assumption that rules out additional sources of information, in particular from asset values or nominal interest rates. A richer model that embeds these additional information sources would speed up the learning process and mitigate the small but persistent long-run effects without changing the models' significant short run implications.

Although these results represent a challenge for models with heterogeneous information as the only source of nominal non-neutrality, they also point to directions for future research: the large adjustment of prices on impact, due to the presence of firm specific shocks, is due in part to an assumption of complete flexibility of prices, along with the premise that firms update based on noisy information in real time. It is plausible to think that adding small amounts of nominal rigidities or menu costs would reduce the adjustment in the short run, and might also lead to more important real effects at medium to long horizons, i.e. once the emphasis of the firms' filtering problem shifts from the transient nature of firm-specific shocks to the problem of inference about aggregate variables.

Future work will have to explore these issues, which generate substantive as well as technical challenges. In particular, the infinite regress problem and the associated aggregation issues, which we circumvented with the finite-horizon assumption, along with the linearity of our flexible adjustment model, will come back in full force in the non-linear pricing model that would emerge from a combination of menu costs with heterogeneous information.[19]

[19] See Gorodnichenko, (2008), for an attempt at tackling these issues.

References

Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (5): 947–985.

Burstein, Ariel T., and Christian Hellwig. 2007. "Prices and Market Shares in a Menu Cost Model." NBER Working Paper No. 13455.

Caplin, Andrew, and Daniel Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102 (November): 703–25.

De Gregorioa, Jose, Alberto Giovanninic, and Holger C. Wolfe. 1994. "International evidence on tradables and nontradables inflation." *European Economic Review* 38 (6): 1225–1244.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Tavares Rebelo. 2008. "Reference Prices and Nominal Rigidities." NBER Working Paper No. 13829.

Golosov, Mikhail, and Robert E. Lucas Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (2): 171–199.

Gorodnichenko, Yuriy. 2008. "Endogenous Information, Menu Costs and Inflation Persistence."

Hellwig, Christian. 2002. "Public Announcements, Adjustment Delays and the Business Cycle."

———. 2008a. "Heterogeneous Information and Business Cycle Fluctuations." work in progress, presented at the 2008 ASSA annual meetings (New Orleans).

———. 2008b. "Monetary Business Cycles (Imperfect Information)." In *The New Palgrave Dictionary of Economics*, edited by Steven N. Durlauf and Lawrence E. Blume. London: Palgrave-Macmillan.

Hellwig, Christian, and Laura Veldkamp. 2008. "Knowing What Others Know: Coordination Motives in Information Acquisition."

Klenow, Peter J., and Oleksiy Kryvstov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123 (3): 863–904.

Lorenzoni, Guido. 2008. "A Theory of Demand Shocks." NBER Working Paper No. 12477.

Lucas, Robert Jr. 1972. "Expectations and the neutrality of money." *Journal of Economic Theory* 4 (2): 103–124 (April).

Mackowiak, Bartosz Adam, and Mirko Wiederholt. 2008. "Optimal Sticky Prices Under Rational Inattention."

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.

Midrigan, Virgiliu. 2007. "Menu Costs, Multi-Product Firms and Aggregate Fluctuations."

Nakamura, Emi, and Jon Steinsson. 2008. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415–1464.

Phelps, Edmund S. 1970. "Introduction: The New Microeconomics in Employment and Inflation Theory." In *Microeconomic Foundations of Employment and Inflation Theory*, edited by Edmund et al. Phelps. New York: Norton.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (5): 665–690.

Townsend, Robert M. 1983. "Forecasting the Forecasts of Others." *Journal of Political Economy* 91 (4): 546.

Woodford, Michael D. 2001. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, edited by Phillipe Aghion et al.

———. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

7 Appendix: Proofs

Proof of Propositions 1 and 2:. Follows immediately from the characterization of k in (3). ■

Proof of Lemma 1:. Given the definition of Γ and Δ, the signal vector can be rewritten as

$$\begin{pmatrix} X_{it} - \hat{X}_{it} \\ \Omega_{it} - \hat{\Omega}_{it} \end{pmatrix} = \Gamma U_t + \Delta \begin{pmatrix} V_{it}^1 \\ \tilde{V}_{it}^1 \\ V_{it}^2 \\ \tilde{V}_{it}^2 \end{pmatrix}$$

Therefore, the vector $\left(U_t', V_{it}^{1\prime}, \tilde{V}_{it}^{1\prime}, V_{it}^2, \tilde{V}_{it}^{2\prime}; X_{it}' - \hat{X}_{it}', \Omega_{it}' - \hat{\Omega}_{it}'\right)'$ is normally distributed with mean zero and variance-covariance matrix:

$$\begin{bmatrix} \multicolumn{2}{c}{\multirow{2}{*}{I}} & \Gamma' \\ & & \Delta' \\ \Gamma & \Delta & \Gamma\Gamma' + \Delta\Delta' \end{bmatrix}$$

and $\left(U_t', V_{it}^{1\prime}, \tilde{V}_{it}^{1\prime}, V_{it}^{2}, \tilde{V}_{it}^{2\prime}\right)'$, conditional on $\left(X_{it} - \hat{X}_{it}, \Omega_{it} - \hat{\Omega}_{it}\right)$ is normally distributed with mean

$$\begin{pmatrix} \Gamma' \\ \Delta' \end{pmatrix} \left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1} \begin{pmatrix} X_{it} - \hat{X}_{it} \\ \Omega_{it} - \hat{\Omega}_{it} \end{pmatrix}.$$

Averaging over i and using the above characterization of signals, we find that average expectations of $\left(U_t', V_{it}^{1\prime}, \tilde{V}_{it}^{1\prime}, V_{it}^{2}, \tilde{V}_{it}^{2\prime}\right)'$ are

$$\begin{pmatrix} \Gamma' \\ \Delta' \end{pmatrix} \left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1} \Gamma U_t$$

From this, the characterization follows immediately. Moreover, the posterior variance-covariance matrix of $\left(U_t', V_{it}^{1\prime}, \tilde{V}_{it}^{1\prime}, V_{it}^{2}, \tilde{V}_{it}^{2\prime}\right)'$ is

$$I - \begin{pmatrix} \Gamma' \\ \Delta' \end{pmatrix} \left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1} \begin{pmatrix} \Gamma & \Delta \end{pmatrix},$$

so that $I - \Gamma' \left(\Gamma\Gamma' + \Delta\Delta'\right)^{-1} \Gamma$ denotes the posterior variance covariance matrix of U_t. ■

Proof of Proposition 3:. When $\tilde{\sigma}_b^2 = \tilde{\sigma}_z^2 = 0$, ϕ' simplifies to

$$\phi = \frac{\kappa\psi^{-1}\rho_b + \rho_z}{1 + \theta\kappa}\mathbf{1} + r\rho_b\hat{\phi}.$$

Using the fact that $\hat{\phi} = (0, \phi_1, ..., \phi_{T-1})' = \Lambda\phi$, where

$$\Lambda = \begin{pmatrix} 0 & 0 \\ I_{T-1} & 0 \end{pmatrix},$$

we have

$$\phi = \frac{\kappa\psi^{-1}\rho_b + \rho_z}{1 + \theta\kappa} \left[I - r\rho_b\Lambda\right]^{-1} \cdot \mathbf{1}$$

The result then follows from

$$[I - r\rho_b\Lambda]^{-1} = \begin{pmatrix} 1 & 0 & \cdot & \cdot & 0 & 0 \\ -r\rho_b & 1 & \cdot & \cdot & \cdot & 0 \\ 0 & -r\rho_b & \cdot & \cdot & \cdot & \cdot \\ \cdot & 0 & \cdot & 1 & 0 & \cdot \\ 0 & \cdot & \cdot & -r\rho_b & 1 & 0 \\ 0 & 0 & \cdot & 0 & -r\rho_b & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & 0 & \cdot & \cdot & 0 & 0 \\ r\rho_b & 1 & \cdot & \cdot & \cdot & 0 \\ (r\rho_b)^2 & r\rho_b & \cdot & \cdot & \cdot & \cdot \\ \cdot & (r\rho_b)^2 & \cdot & 1 & 0 & \cdot \\ (r\rho_b)^{T-2} & \cdot & \cdot & r\rho_b & 1 & 0 \\ (r\rho_b)^{T-1} & (r\rho_b)^{T-2} & \cdot & (r\rho_b)^2 & r\rho_b & 1 \end{pmatrix}$$

■