Incomplete Information, Higher-Order Beliefs and the Inertia of Prices in the Calvo Model

George-Marios Angeletos — MIT and NBER

Jennifer La'O — MIT

November 22, 2008

Abstract

This paper investigates how incomplete information interacts with sticky prices in determining the response of prices to nominal shocks. Our baseline model is a variant of the Calvo model in which firms observe noisy private signals about the underlying nominal shocks. The response of prices is then pinned down by three parameters: the precision of information about these shocks; the frequency of price adjustment; and the degree of strategic complementarity in pricing decisions. This result synthesizes the broader lessons of the pertinent literature. We next proceed to highlight why these lessons provide only a partial picture: in general, the precision of information about nominal shocks does not necessarily pin down the response of higher-order beliefs to these shocks. By implication, one cannot calibrate the degree of price inertia without additional information about the behavior of higher-order beliefs. We highlight the crucial, and distinct, role of higher-order beliefs with three extensions of our baseline model. First, we show how uncertainty about the precision of the information of other firms can impact the response of prices without impacting the response of the firms' expectations about these shocks. Second, we show how heterogeneous priors about the information of other firms can facilitate significant price inertia even if both the precision of information and the frequency of price adjustment are arbitrarily high. Finally, we show how independent variation in higher-order beliefs can generate fluctuations that resemble "cost-push" shocks.

1 Introduction

How much, and how quickly, do prices respond to nominal shocks? This is one of the most fundamental questions in macroeconomics: it is key to understanding both the sources of the business cycle and the power of monetary policy to control real economic activity.

To address this question, one strand of the literature has focused on menu costs and other frictions in adjusting prices; this includes both time-dependent and state-dependent adjustment models. Price rigidities are then identified as the key force behind price inertia. Another strand of the literature has focused on informational frictions; this strand highlights that firms may fail to adjust their price to nominal shocks, not because it is costly or impossible to do so, but rather because they have only imperfect information about these shocks. The precision of available information about these shocks then emerges as the key determinant of the response of prices to nominal shocks.

The starting point of this paper is a bridge between these two approaches. In particular, our baseline framework is a hybrid of Calvo (1983) and Woodford (2003): on the one hand, firms can adjust prices only infrequently, as in Calvo; on the other hand, firms observe the underlying nominal shocks only with noise, as in Woodford.[1]

Within this model, the response of prices to nominal shocks—and hence also the real impact of these shocks—is characterized by the interaction of three key parameters: the precision of available information about these shocks; the frequency of price adjustment; and the degree of strategic complementarity in pricing decisions. In particular, the combination of sticky prices and strategic complementarity implies that the incompleteness of information can have lasting effects on inflation and real output even if the shocks become commonly known very quickly. This is because firms that have full information about the shock at the time they set prices will find it optimal to adjust only partly to the extent that other firms had only incomplete information at the time they had set their prices. Moreover, incomplete information can help make inflation peak after real output, which seems consistent with available evidence based on structural VARs.

These findings synthesize, and marginally extend, various lessons from the pertinent literature

[1] Although the available evidence appears to favor state-dependent models over time-dependent models, we employ the Calvo model for two reasons. The one is mere convenience: the Calvo model remains tractable even after we add incomplete information, whereas state-dependent models are often much less tractable even with complete information. The other is our belief that the distinction between time- and state-dependent models need not be crucial for the particular insights of this paper.

with regard to how frictions in price adjustment and noisy information about the underlying nominal shocks impact the response of prices to these shocks. Although this synthesis has some value on its own, this is not the main contribution of the paper. Rather, we use the synthesized framework in order to highlight a different lesson, one that regards the dynamic behavior of higher-order beliefs.

The precision of available information about the underlying nominal shock identifies how fast the firms' forecasts of the shock respond to the true shock: the more precise the available information, the faster the firms' forecasts converge to the truth. However, this need not also identify how fast the forecasts of the forecasts of others may respond. In other words, the precision of available information pins down the response of first-order beliefs, but not necessarily the response of higher-order beliefs. But when prices are strategic complements, the response of the price level to the underlying shock depends heavily on the response of higher-order beliefs. It follows that neither the precision of available information nor the degree of price rigidity suffice for calibrating the degree of price inertia at the macro level: rather, one also needs appropriate information about the behavior of higher-order beliefs.

To understand this property, it is useful to abstract for a moment from sticky prices. Assume, in particular, that all firms can adjust their prices in any given period and that the prices they set are given by the following simple pricing rule:

$$p_i = (1-\alpha)\mathbb{E}_i\theta + \alpha\mathbb{E}_i p$$

where p_i is the price set by firm i, θ is nominal demand, p is the aggregate price level, $\alpha \in (0,1)$ is the degree of strategic complementarity in pricing decisions, and $\mathbb{E}_i$ denotes the expectation conditional on the information of firm i. Aggregating this condition and iterating over the expectations of the price level, we infer that the aggregate price level must satisfy the following condition:

$$p = (1-\alpha)\left(\bar{E}^1 + \alpha\bar{E}^2 + \alpha^2\bar{E}^3 + \dots\right),$$

where $\bar{E}^k$ denotes the k-order average forecast of θ. The k-order average forecasts are defined recursively as follows: $\bar{E}^1$ is the cross-sectional mean of the firms' forecasts of the underlying nominal shock; $\bar{E}^2$ is the cross-sectional mean of the firms' forecasts of $\bar{E}^1$; and so on.[2] It then follows that understanding the response of price level p to the nominal shock θ boils down to understanding the response of the sequence of different orders of beliefs, $\{\bar{E}^k\}_k$, to that shock.

[2]Formally, $\bar{E}^1 = \mathbb{E}[\mathbb{E}_i\theta|\theta]$ and $\bar{E}^k = \mathbb{E}[\mathbb{E}_i E^{k-1}|\theta]$ for all $k \geq 2$.

We highlight the crucial and distinct role of higher-order beliefs within three variants of our baseline model. All three extensions retain the combination of infrequent price adjustment a la Calvo and noisy information a la Woodford, but differentiate in the specification of higher-order beliefs.

Our interest in highlighting this distinction is motivated by the following. When information is complete or at least common across all firms, then $\bar{E}^k = \bar{E}^1$ for all k. It then follows that the response of prices to nominal shocks depends merely on the response of first-order beliefs to those shocks, which in turn is merely a function of the precision of available information. When, instead, information is incomplete and dispersed, higher-order beliefs need not coincide with first-order beliefs. However, in the standard Gaussian example used in most of the literature, the sensitivity of higher-order beliefs to the shock is tightly connected to that of first-order beliefs: higher-order beliefs can be less sensitive to the underlying shock only if the precision of information about the shock is lower, in which case first-order beliefs are also less sensitive. It follows that in the standard Gaussian example the precision of available information about the underlying nominal shock remains the key determinant of the response of the price level to nominal shocks, much alike the case with only common uncertainty.[3] However, once one goes away from the standard Gaussian example, this tight connection between first- and higher-order beliefs may well break.

We show that this is indeed what happens in our three variants of our baseline model. In the first variant, firms face uncertainty, not only about the size of the aggregate nominal shock, but also about the precision of the signals that *other* firms receive about this shock. This extension helps isolate the role of higher-order beliefs or, equivalently, the role of strategic uncertainty: we show how this additional source of uncertainty about the distribution of precisions can impact the response of higher-order beliefs to the underlying shocks, and thereby the response of prices, without affecting the response of first-order beliefs

In the second variant, we let firms hold heterogeneous priors about the stochastic properties of the signals that other firms receive. In this economy, firms expect the beliefs of others to adjust more slowly to the underlying shocks than their own beliefs. They thus behave in equilibrium as if they lived in a economy where all other firms had more imprecise information that what they do. This in turn helps rationalize why equilibrium prices may adjust very slowly to the underlying nominal

[3]The only differences are (i) that the degree of complementarity now matters, because it controls how much the price level depends on higher-order beliefs; and (ii) the decomposition of information between private and public also matters, because higher-order beliefs are more anchored towards public information.

shocks even if the frequency of price adjustment is arbitrarily high and each firm has arbitrarily precise information about the underlying nominal shock.

The aforementioned two variants focus on how higher-order beliefs impact the propagation of nominal shocks in the economy. In the third and final variant, we show how higher-order beliefs can be the source of fluctuations in the economy—how they can themselves be one of the "structural" shocks. In particular, we show how variation in higher-order beliefs that is orthogonal to either the underlying nominal shocks or the firms information about these shocks can generate fluctuations in inflation and real output that resemble those generated by"cost-push" shocks.

Combined, these findings point out that a macroeconomist who wishes to quantify the response of the economy to the underlying structural shocks, or even to identify what are these structural shocks in the first place, may need appropriate information, not only about the degree of price rigidity and the the firms' information (beliefs) about these shocks, but also about the stochastic properties of their beliefs about the beliefs of others.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 studies our baseline model, which introduces incomplete information to the Calvo model. Section 4 studies the variant where firms face uncertainty about the precisions of the signals of other firms. Section 5 studies the variant that allows for heterogeneous priors. Section 6 turns to cost-push shocks. Section 7 concludes. All proofs are in the Appendix.

2 Related literature

The macroeconomics literature on informational frictions has a long history, going back to Phelps (1971), Lucas (1972), Barro (1976), and Townsend (1983). Recently, this literature has been revived by Mankiw and Reis (2002), Morris and Shin (2002), Sims (2003), Woodford (2003), and other subsequent work.[4] This paper contributes to this literature in two ways: first, by studying the interaction of incomplete information with price rigidities; second, and most importantly, by furthering our understanding of the role of higher-order beliefs for the response of prices at the macro level.

Our paper is highly complementary to the papers by Woodford (2003) and Morris and Shin

[4] See, among others, Amato and Shin (2006), Angeletos and La'O (2008a), Angeletos and Pavan (2007), Bacchetta and Wincoop (2005), Collard and Dellas (2005), Hellwig (2005), Lorenzoni (2008), Mackowiak and Wiederholt (2007), and Reis (2006).

(2002, 2006). These papers document how higher-order beliefs may respond less to information about the underlying shocks than first-order beliefs, simply because they are more anchored to the common prior. However, by restricting attention to a very specific information structure, they have also restricted attention to settings where the response of higher-order beliefs is tightly tied to the response of first-order beliefs: in their settings, the response of higher-order beliefs is a monotone transformation of the response of first-order beliefs, thus precluding any *independent* role for higher-order beliefs. Our paper, instead, highlights that, whereas the response of first-order beliefs to the underlying nominals shocks is pinned down solely by the level of uncertainty about these shocks, the response of higher-order beliefs depends also on other sources of uncertainty, such as uncertainty about the precision of others' information. Furthermore, it shows how with heterogeneous priors it is possible that higher-order beliefs respond very little, or even not at all, to the underlying shocks even if all agents are nearly perfectly informed about these shocks (in which case first-order beliefs respond nearly one-to-one with the shock).

Also related is Klenow and Willis (2007). This paper considers a model that combines menu costs with sticky information a la Mankiw and Reis (2002), focusing on how significant price inertia at the macro level can be consistent with significant price flexibility at the micro level to the extent that firms update their information about the underlying aggregate nominal shocks at a sufficiently slower rate than their information about their idiosyncratic shocks.[5] In contrast, our paper shows that significant price inertia at the macro level is consistent, not only with significant price flexibility at the micro level, but also with fast learning about the underlying nominal shocks.

Finally, also related are Scheinkman and Xiong (2003) and Angeletos and La'O (2008b). The first paper shows how heterogeneous priors can help explain speculative asset-price bubbles. The second paper shows how they can help sustain sunspot-like fluctuations—i.e., variation in aggregate employment, output and consumption that cannot be explained by variation in *either* the underlying economic fundamentals or the agents' information about these fundamentals—in an otherwise standard real-business-cycle model with a unique equilibrium. The present paper complements this prior work by illustrating another dimension in which heterogeneous priors can be useful modeling devices: they can help rationalize significant inertia in the response of prices to nominal shocks.

[5]Building upon Sim's (2001) theory of rational inattention, Mackowiak and Wiederholt (2007) provide a simple intuition for why this may be the case: if allocating attention to different pieces of information is costly (or subject to a capacity constraint), it may be natural to expect firms to allocate less attention to macroeconomic shocks than to idiosyncratic shocks simply because the former are much less volatile than the latter.

3 The Calvo Model with Incomplete Information

In this section we consider a variant of the Calvo model that allows firms to have dispersed private information about aggregate nominal demand.

Households and firms. The economy is populated by a representative household and a continuum of firms that produce differentiated commodities. Firms are indexed by $i \in [0,1]$. Time is discrete, indexed by $t \in \{0,1,2,...\}$. There is no capital, so that there is no saving in equilibrium. Along with the fact that there is a representative household, we can also abstract from asset trading.

The per-period utility of the household is given by

$$U_t = \log C_t - N_t,$$

where N_t is the labor supplied by the household,

$$C_t = \left[\int C_{i,t}^{\frac{\eta-1}{\rho}} di\right]^{\frac{\eta}{\eta-1}}$$

is the familiar CES aggregator, $C_{i,t}$ is the consumption of the commodity produced by firm i, and $\eta > 0$ is the elasticity of substitution across commodities. As usual, this specification implies that the demand for the commodity of firm i is given by

$$C_{i,t} = \left(\frac{P_{i,t}}{P_t}\right)^{-\eta} C_t,$$

where $P_t \equiv \left[\int P_{i,t}^{\eta-1} di\right]^{\frac{1}{\eta-1}}$ is the aggregate price index. The output of firm i, on the other hand, is given by

$$Y_{i,t} = A_{i,t} L_{i,t}^{\epsilon},$$

where $A_{i,t}$ is the idiosyncratic productivity shock and $\epsilon \in (0,1)$ parameterizes the degree of diminishing returns. In equilibrium, $Y_{i,t} = C_{i,t}$ for all firms and aggregate output is simply $Y_t = C_t$. Finally, households face a binding cash-in-advance constraint of the form

$$P_t C_t = \Theta_t,$$

where Θ_t denotes aggregate nominal demand (aggregate nominal GDP). The latter is assumed to be exogenous and defines the "fundamentals" for our model.

In what follows, we use lower-case variables for the logarithms of the corresponding upper-case variables: $\theta_t \equiv \log \Theta_t$, $y_t = \log Y_t$, $p_t = \log P_t$, and so on. We also assume that all exogenous shocks are log-normally distributed, which guarantees that the equilibrium admits an exact log-linear solution.

Shocks and information. Aggregate nominal demand is assumed to follow a random walk:

$$\theta_t = \theta_{t-1} + v_t$$

where $v_t \sim \mathcal{N}(0, \sigma_\theta^2)$ is white noise. Each period has two stages, a morning and an evening. Information about the current level of nominal demand is imperfect during the morning but perfect during the evening.[6] The information that a firm has about θ_t during the morning is summarized in a Gaussian private signal of the following form:

$$x_{i,t} = \theta_t + \varepsilon_{i,t},$$

where $\varepsilon_{i,t} \sim \mathcal{N}(0, \sigma_x^2)$ is purely idiosyncratic noise (i.i.d. across firms). Pricing choices (for the firms that have the option to set prices) are made in the morning, while information about θ_t is imperfect; employment and consumption choices are made in the evening, once θ_t has been publicly revealed. Finally, we assume that the idiosyncratic productivity shock $a_{i,t}$ follows a random walk and that it is known to the firm from the beginning of the period.

The information set of firm i at the morning of date t is therefore given by

$$\mathcal{I}_{i,t} = \{\theta_{t-1}, \theta_{t-2}, \ldots; x_{i,t}, x_{i,t-1}, x_{i,t-2}, \ldots; a_{i,t}, a_{i,t-1}, a_{i,t-2}, \ldots\} = \mathcal{I}_{i,t-1} \cup \{\theta_{t-1}, x_{i,t}, a_{i,t}\},$$

while its information set at the evening of the same date is given by $\mathcal{I}_{i,t} \cup \{\theta_t\}$. (We will see in a moment that, in equilibrium, p_{t-1} and y_{t-1} are functions of $\theta_{t-1}, \theta_{t-2}, \ldots$; it would thus make no difference if we had included past values of the price level and real GDP in the information set of the firm. Similarly, $y_{i,t}$ is a function of $\mathcal{I}_{i,t}$; it would thus make no differed if we had included the realized level of a firm's demand into its evening information set.)

Price-setting behavior. The "target" price of firm i in period t—i.e., the price the firm would have set in the morning of that period if prices had been flexible and information had been complete—is given by

$$p_{i,t}^* = \mu + mc_{i,t}$$

where $\mu \equiv \frac{\eta}{\eta-1}$ is the monopolistic mark-up and $mc_{i,t}$ is the nominal marginal cost the firm faces in the evening of period t. The latter is given by

$$mc_{i,t} = w_t + \frac{1-\epsilon}{\epsilon} y_{i,t} - \frac{1}{\epsilon} a_{i,t}$$

[6] This morning-evening structure is borrowed from Angeletos and La'O (2008), which study a class of micro-founded real and monetary business-cycle economies with imperfect competition and incomplete information. See also Hellwig (2005), and Lorenzoni (2008) for related contributions.

where w_t is the nominal wage rate in period t. From the representative household's optimality condition for work,

$$w_t - p_t = c_t.$$

From the consumer's demand,

$$c_{i,t} - c_t = -\eta(p_{i,t} - p_t).$$

From market clearing, $c_{i,t} = y_{i,t}$ and $c_t = y_t$. Finally, from the cash-in-advance constraint, aggregate real output is given by

$$y_t = \theta_t - p_t.$$

Combining the aforementioned conditions, we conclude that the target price of firm i in period t is given by

$$p^*_{i,t} = (1-\alpha)\theta_t + \alpha p_t + \xi_{i,t} \tag{1}$$

where

$$\alpha \equiv 1 - \frac{1}{\epsilon + (1-\epsilon)\eta} < 1$$

defines the degree of strategic complementarity in prices and $\xi_{i,t} \equiv \frac{\epsilon}{\epsilon+(1-\epsilon)\eta}\mu - \frac{1}{\epsilon+(1-\epsilon)\eta}a_{i,t}$ is simply a linear transformation of the idiosyncratic productivity shock. We henceforth restrict $\alpha > 0$ and normalize the mean of the idiosyncratic productivity shock so that the cross-sectional mean of $\xi_{i,t}$ is zero.

If prices had been flexible and θ_t had been known in the morning of period t, the firm would set $p_{i,t} = p^*_{i,t}$ in all periods and states. However, following Calvo (1983), we assume that a firm may change its price only with probability $1-\lambda$ during any given period, where $\lambda \in (0,1)$. It then follows that, in the event that a firm changes its price, the price it chooses is equal (up to a constant) to a weighted average of all future expected target prices:

$$p_{i,t} = \mathbb{E}_{i,t}\left[(1-\beta\lambda)\sum_{j=0}^{\infty}(\beta\lambda)^j\, p^*_{t+j}\right] \tag{2}$$

where $\beta \in (0,1)$ is the discount factor, λ is the probability that the firm won't have the option to adjust its price (a measure of how sticky prices are), and $\mathbb{E}_{i,t}$ is the expectation conditional on the information set of firm i in period t.[7]

[7] To be precise, condition (2) should have been written as $p_{i,t} = const + \mathbb{E}_{i,t}\left[(1-\beta\lambda)\sum_{j=0}^{\infty}(\beta\lambda)^j\, p^*_{t+j}\right]$, where *const* is a an endogenous quantity that involves second-order moments and that emerges due to risk aversion. However, because these second-order moments are invariant with all the shocks and the information of the firms, it is without any loss of generality to ignore it.

Combining conditions (1) and (2), we conclude that the price set by any firm that gets the chance to adjust its price in period t is given by

$$p_{i,t} = (1-\beta\lambda)\sum_{j=0}^{\infty}(\beta\lambda)^j\left[(1-\alpha)\,\mathbb{E}_{i,t}\theta_{t+j} + \alpha\mathbb{E}_{i,t}p_{t+j} + \mathbb{E}_{i,t}\xi_{i,t+j}\right] \tag{3}$$

In the remainder of the paper, we treat (3) as it were an exogenously-defined behavioral rule, with the understanding though that this rule is actually fully rationalized in equilibrium.

Equilibrium dynamics. The economy effectively reduces to a dynamic game of incomplete information, with condition (3) representing the best response of the typical firm. The equilibrium notion we adopt is standard PBE.[8] Because of the linearity of the best-response condition (3) and the Gaussian specification of the information structure, it is a safe guess that the equilibrium strategy will have a linear form. We thus conjecture the existence of equilibria in which the price set by a firm in period t is a linear function of $(p_{t-1}, \theta_{t-1}, x_{i,t}, \xi_{i,t})$:

$$p_{i,t} = P(p_{t-1}, \theta_{t-1}, x_{i,t}, \xi_{i,t}) = b_1 p_{t-1} + b_2 x_{i,t} + b_3\theta_{t-1} + \xi_{i,t} \tag{4}$$

for some coefficients b_1, b_2, b_3. This particular guess is justified by the following reasoning: we expect p_{t-1} to matter because of a fraction of firms cannot adjust prices; $x_{i,t}$ because it conveys information about the current nominal shock θ_t; θ_{t-1} because it is the prior about θ_t; and $\xi_{i,t}$ for obvious reasons.

Given this guess, and given the fact that only a randomly selected fraction $1-\lambda$ of firms can adjust prices in any given period, we infer that the aggregate price level must satisfy

$$p_t = \lambda p_{t-1} + (1-\lambda)\int\int P(p_{t-1}, \theta_{t-1}; x, \xi)\,dF_t(x)\,dG_t(\xi)$$

where F_t denotes the cross-sectional distribution of the private signals (conditional on the current shock θ_t) and G_t denotes the cross-sectional distribution of the idiosyncratic shocks. Given that P is linear, that the cross-sectional average of x_t is θ_t, and that cross-sectional average of ξ_t is 0, we can re-write the above as $p_t = \lambda p_{t-1} + (1-\lambda)\,P(p_{t-1}, \theta_{t-1}; \theta_t, 0)$, or equivalently as

$$p_t = c_1 p_{t-1} + c_2\theta_t + c_3\theta_{t-1} \tag{5}$$

[8] We can safely ignore out-of-equilibrium paths by assuming that firms observe (at most) the cross-section distribution of prices, not the specific prices set by specific firm. The fact that each firm is infinitesimal then guarantees that no firm's atomistic deviation can cause other firms to detect a deviation of the equilibrium path.

where

$$c_1 = \lambda + (1-\lambda)b_1, \quad c_2 = (1-\lambda)b_2 \quad c_3 = (1-\lambda)b_3. \tag{6}$$

Next, note that we can rewrite (3) in recursive form as

$$p_{i,t} = (1-\beta\lambda)\left[(1-\alpha)\,\mathbb{E}_{i,t}\theta_t + \alpha\mathbb{E}_{i,t}p_t + \xi_{i,t}\right] + (\beta\lambda)\mathbb{E}_{i,t}p_{i,t+1}.$$

Using (4) and (5) into the right-hand side of the above condition, we infer that the price must satisfy

$$p_{i,t} = (1-\beta\lambda)\left[(1-\alpha)\,\mathbb{E}_{i,t}\theta_t + \alpha\mathbb{E}_{i,t}p_t + \xi_{i,t}\right] + (\beta\lambda)\left[b_1\mathbb{E}_{i,t}p_t + b_2\mathbb{E}_{i,t}\theta_t + b_3\mathbb{E}_{i,t}\theta_{t+1} + \mathbb{E}_{i,t}\xi_{i,t+1}\right]. \tag{7}$$

Next, note that

$$\mathbb{E}_{i,t}\theta_{t+1} = \mathbb{E}_{i,t}\theta_t = \frac{\kappa_x}{\kappa_x+\kappa_\theta}x_{it} + \frac{\kappa_x}{\kappa_x+\kappa_\theta}\theta_{t-1} \qquad \text{and} \qquad \mathbb{E}_{i,t}\xi_{i,t+1} = \xi_{i,t},$$

where $\kappa_x \equiv \sigma_x^{-2}$ is the precision of the firms' signals and $\kappa_\theta \equiv \sigma_\theta^{-2}$ is the precision of the common prior about the innovation in θ. Using these facts, and substituting p_t from (5) into (7), we can rewrite the left-hand side of (7) as a linear function of p_{t-1}, $x_{i,t}$, θ_{t-1}, and $\xi_{i,t}$. For this to coincide with our conjecture in (4), it is necessary and sufficient that the coefficients (b_1, b_2, b_3) solve the following system:

$$\begin{aligned}
b_1 &= (1-\beta\lambda)\alpha c_1 + (\beta\lambda)b_1c_1 \\
b_2 &= \left[(1-\beta\lambda)(1-\alpha+\alpha c_2) + (\beta\lambda)(b_1c_2+b_2+b_3)\right]\frac{\kappa_x}{\kappa_x+\kappa_\theta} \\
b_3 &= \left[(1-\beta\lambda)(1-\alpha+\alpha c_2) + (\beta\lambda)(b_1c_2+b_2+b_3)\right]\frac{\kappa_\theta}{\kappa_x+\kappa_\theta} \\
&\quad +(1-\beta\lambda)\alpha c_3 + (\beta\lambda)b_1c_3
\end{aligned} \tag{8}$$

We conclude that an equilibrium is pinned down by the joint solution of (6) and (8).

Combining the conditions for c_1 and b_1, we get that c_1 must solve the following equation:

$$c_1 = \lambda + (1-\lambda)\left(\frac{1-\beta\lambda}{1-\beta\lambda c_1}\right)\alpha c_1.$$

This equation admits two solutions: one with $c_1 > 1$ and another with $c_1 \in (\lambda, 1)$. We ignore the former solution because it leads to explosive price paths and henceforth limit attention to the latter solution. Note that this solution is independent of the information structure; indeed, the coefficient c_1, which identifies the endogenous persistence in the price level, coincides with the one in the standard (complete-information) Calvo model.

Given this solution for c_1, the remaining conditions define a linear system that admits a unique solution for the remainder of the coefficients. It is straightforward to check that the solution satisfies

$$c_1 + c_2 + c_3 = 1,$$

which simply means that the price process is homogenous of degree one in the level of nominal demand. Furthermore,

$$c_2 = \frac{\lambda(1 - c_1)}{\lambda + c_1 \frac{\kappa_\theta}{\kappa_x}}$$

which identifies the sensitivity of the price level to the current innovation in nominal demand as an increasing function of the precision of available information. We therefore reach the following characterization of the equilibrium.

Proposition 1. *(i) There exists an equilibrium in which the pricing strategy of a firm is given by*

$$p_{i,t} = b_1 p_{t-1} + b_2 x_{i,t} + b_3 \theta_{t-1}$$

and the aggregate price level is given by

$$p_t = c_1 p_{t-1} + c_2 \theta_t + c_3 \theta_{t-1}$$

for some positive coefficients (b_1, b_2, b_3) *and* (c_1, c_2, c_3).

(ii) The equilibrium values of the coefficients (c_1, c_2, c_3) *satisfy the following properties:* c_1 *is increasing in* λ*, increasing in* α*, and invariant to* κ_x/κ_θ*;* c_2 *is non-monotone in* λ*, decreasing in* α*, and increasing in* κ_x/κ_θ*;* c_3 *is non-monotone in* α*, and decreasing in* κ_x/κ_θ*.*

The comparative statics described in the above proposition are illustrated in Figures 1, 2 and 3. The baseline parameterization we use for these figures, as well as for the impulse responses that we report later on, is as follows. We identify the length of a period with one year; this seems a good benchmark for how long it takes for macro data to become publicly known. We accordingly set $\beta = .95$ (which corresponds to a discount rate of about 1% per quarter), $\lambda = (2/3)^4$ (which corresponds to a third of firms changing prices every quarter), and $\alpha = .85$; these values are broadly consistent with available empirical estimates and with standard calibrations of the Calvo model. In want of any obvious estimate of the precision of information, we set $\kappa_x/\kappa_v = 1$; this means that the variance of the forecast error about the current shock of the typical firm is one half the variance of the innovation in the shock itself.

Figure 1 plots the coefficients c_1, c_2, and c_3 as functions of the Calvo parameter λ, the probability the firm does *not* revise its price in a given period. We observe that $c1$ is increasing in λ, $c3$ is decreasing in λ, and $c2$ is non-monotonic in λ (it increases for low values but decreases for high values). Figure 2 plots these coefficients as functions of α, the degree of strategic complementarity in pricing decisions. We observe that c_1 is an increasing function in α, c_2 is a decreasing function in α, and c_3 is non-monotonic in α. Finally, Figure 3 plots these coefficients as functions of κ_x/κ_θ, the ratio of the precision of private signals to the precision of the prior. We see that c_2 is increasing in this ration, while c_3 is decreasing and c_1 is invariant.

The comparative statics described above are a hybrid of the results found in sticky-price Calvo models and in the incomplete information literature. As in the standard Calvo model, the aggregate price level is persistent due to the fact that some firms cannot adjust prices. In our model, the coefficient which characterizes the persistence of the aggregate price process is c_1. We find that this coefficient is unaffected by the incompleteness of information. In this sense, the persistence of prices in our model is the same as in the standard Calvo model and depends only the rate at which firms may change their prices and the degree of strategic complementarity in pricing decisions. The fact that c_1 is increasing in the Calvo parameter λ should then be familiar: it is almost the mechanical implication of the fact that a fraction λ of firms do not adjust prices. The impact of α on c_1 is also familiar: even under complete information, firms who can adjust their price following a monetary shock will find it optimal to stay closer to the past price level the higher the degree of strategic complementarity between them and the firms that cannot adjust (and that are thus stuck to the past price level).

Where the incompleteness of information has a bite is on the coefficients c_2 and c_3, which characterize, respectively, the impact of the current and the past shock on the current price level for any given past price level. To understand how the precision of information affects these coefficients, consider the choice of the price-setting firm. The price chosen by a firm is a linear combination of past prices and past nominal shocks (which are common knowledge among all firms) and the firm's own expectation of current nominal demand (which is unknown in the current period). Aggregating across firms, we find that the aggregate price level is a linear combination of past price levels and past nominal shocks and of the *average* expectation of θ_t. As in any static incomplete information model with Gaussian signals, the agent's own expectation of the fundamental is merely a weighted combination of his private signal and the common prior, which here coincides with θ_{t-1}. If firms

Figure 1

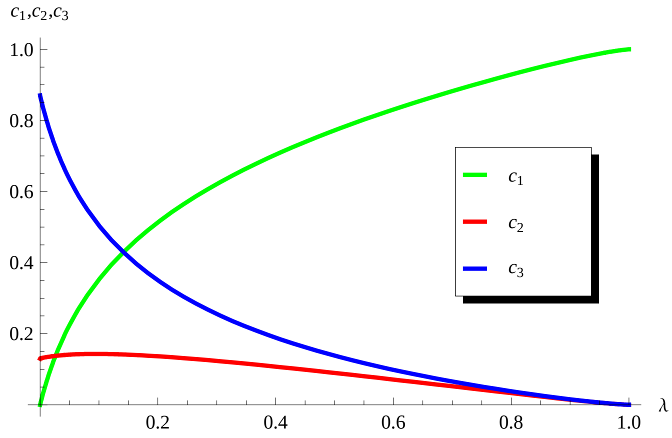

Figure 2

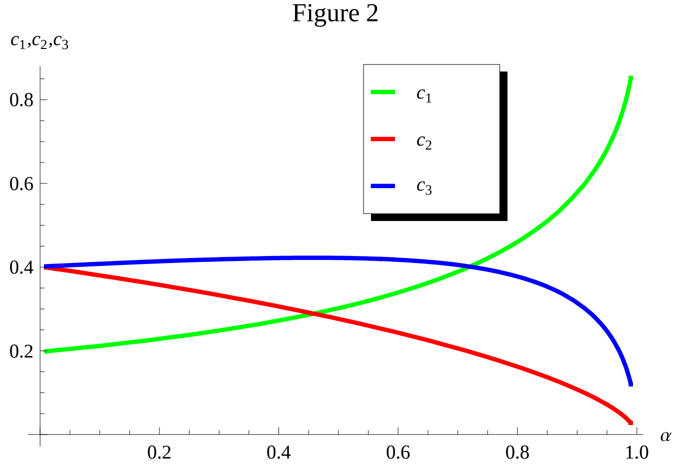

Figure 3

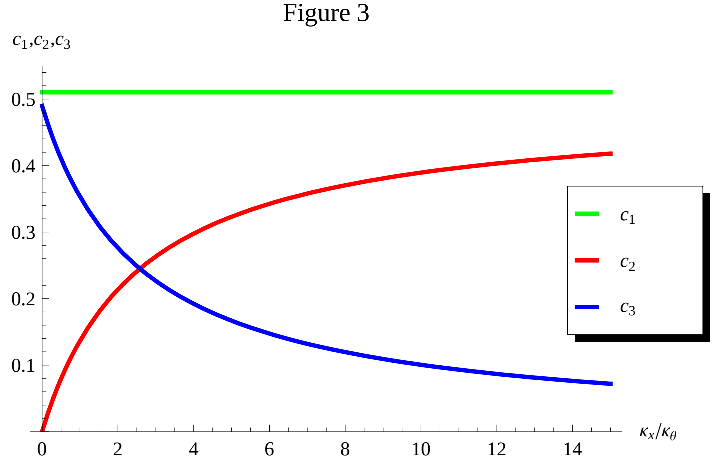

have less precise private information relative to the prior, i.e., lower κ_x/κ_θ, they place less weight on their private signals than on their prior when forming their expectations of θ_t. As a result, the average expectation is less sensitive to the current shock θ_{t-1} and more anchored to the past shock θ_{t-1}. This explains why less precise information (a lower κ_x/κ_θ) implies a lower c_2 and a higher c_3.

Impulse responses. The above analysis highlights how introducing incompleteness of information into the Calvo model dampens the response of prices to the underlying nominal shocks—the precision of information becomes a key parameter for the dynamics of inflation along with the Calvo parameter and the degree of strategic complementarity. To further appreciate this, we now study how the precision of information affects the impulse responses of the inflation rate and real output to an innovation in nominal demand.

Figures 4 and 5 plot these impulse responses. (Inflation in period t is given by $p_t - p_{t-1}$, while real output is $y_t = \theta_t - p_t$.) As before, we identify the period with a year and set $\beta = .95$, $\lambda = (2/3)^4$ and $\alpha = .85$. We then consider three alternative values for the precision of information: $\kappa_x/\kappa_\theta = 1$, which is our baseline; $\kappa_x/\kappa_\theta = \infty$, which corresponds to the extreme of complete information about the current shock (as in the standard Calvo model); and $\kappa_x/\kappa_\theta = 0$, which corresponds to the alternative extreme, that of no information about the current shock other than the prior (i.e., the past shock).

From Figure 4, we see that the incompleteness of information has important effects on inflation dynamics relative to the complete information Calvo model. First, the instantaneous impact effect of a monetary shock on inflation is increasing in κ_x/κ_θ. As the noise in private information increases, prices react less initially to a nominal disturbance. Secondly, as the precision of private information decreases, we observe that second period inflation becomes higher and higher. As the past nominal demand now becomes common knowledge, prices with low sensitivity to the monetary shock last period greatly increase in the second period to reflect this new information. Except for sufficiently high values of κ_x/κ_θ, this is where inflation reaches its peak. Lastly, although the decay rate of inflation is constant after this date, because of the high inflation experienced in the second period, lower κ_x/κ_θ leads to a higher level of inflation for all subsequent periods.

From Figure 5, we then observe that the impact effect of a monetary shock on output is decreasing in the precision of private information. Of course, this is simply the mirror image of what happens to prices. Furthermore, like the impulse responses for inflation, real output is higher for lower levels of $\kappa_\varepsilon/\kappa_v$ for all subsequent periods.

Figure 4

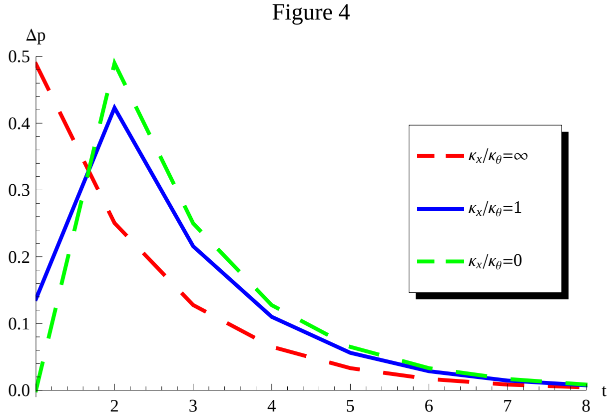

Figure 5

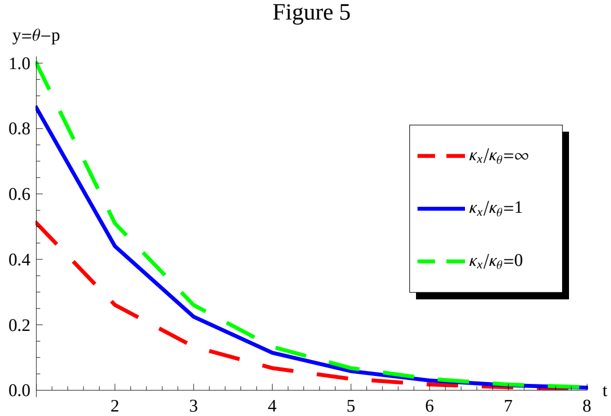

It is interesting here to note that incomplete information has lasting effects on the levels of inflation and real GDP even though the shocks become common knowledge after just one period. This is precisely because of the interaction of incomplete information with price staggering and with strategic complementarity: by the time the shock becomes common knowledge, some firms have already set their price on the basis of incomplete information about the shock; strategic complementarity then guarantees that the firms that now have access to full information will still find it optimal to respond to the shock as if themselves had incomplete information.

Note that, except for high values of κ_x/κ_θ, we observe that the peak of output occurs before the peak in inflation. This is in contrast to the standard Calvo model which predicts strong price increases during the period in which the shock is realized and therefore typically has inflation peaking before output. A similar observation has been made in Woodford (2003), but with two important differences: Woodford (2003) abstracts from price staggering; it also assumes that past shocks and past outcomes never become known, thus appearing to require an implausibly slow degree of learning about the underlying shocks. Here, we show how the empirically appealing property that inflation peaks after real output can be obtained even with quite fast learning, provided one interacts incomplete information with price staggering.

In the present model, the peak of inflation happens at most one period after the innovation in θ. This particular property is an artifact of the assumption that the shock becomes common knowledge exactly one period after the innovation takes place. If we extend the model so that the shock becomes common knowledge after, say, 4 periods, then the peak in inflation can occur as late as 4 periods after the shock. The more general insight is that inflation can start low if agents initially have little information about the innovation and can rise in the early phases of learning, but once agents have accumulated enough information about the shock then inflation will begin to fall. In other words, the dynamics of learning are essential for the dynamics of inflation only as long as the agents remain sufficiently uncertain about the shock; but once the firms have learned enough about the shock, the subsequent dynamics of inflation are determined primarily by the Calvo mechanics.

4 Uncertainty about precisions

The analysis so far has focused on a Gaussian specification for the information structure that is quite standard in the pertinent literature. Under this specification, we show that the response of prices to nominal shocks was determined by three parameters: (i) the degree of price rigidity; (ii) the degree of strategic complementarity; and (iii) the precision of information about the underlying nominal shock. In this section we show that, under a plausible variation of the information structure, knowledge of these parameters need not suffice for calibrating the degree of price inertia. The key insight is that the precision of information about the underlying shock pins down the response of the firms' forecasts of this shock, but not necessarily the response of their forecasts of the forecasts of other firms (i.e., their higher-order beliefs); and what matters for the response of equilibrium prices to the shock is not only the former but also the latter.

Apart from serving as an example for this more general insight, the variant we consider here has its own appeal in that it introduces a plausible source of uncertainty: it allows agents to face uncertainty regarding the precision of information that other firms may have regarding nominal demand.

In particular, we introduce a second aggregate shock, which permits us to capture uncertainty about the average precision of available information in the cross-section of the economy. Let this shock be a binary random variable, $s_t \in \{h, l\}$. This shock is i.i.d. over time, and independent of θ_t, taking each of the two possible values h and l with probability 1/2. Let $\gamma, \kappa_h, \kappa_l$ be scalars, commonly known to all agents, with $1/2 < \gamma < 1$ and $0 < \kappa_l < \kappa_h$. To simplify the notation, we adopt the convention that $-s = l$ when $s = h$ and $-s = h$ when $s = l$. When the realized shock is s, the type of an agent is (x, κ_s) with probability γ and (x, κ_{-s}) with probability $(1-\gamma)$, where x is the particular realization of the private signal about θ that the agent receives and κ is the precision of this signal.

Note that an agent knows his κ, but not the underlying state s. An agent's κ thus serves a double role: it is both the precision of the agent's own information about θ and a noisy signal of the average precision in the cross-section of the economy. Therefore, the key difference from the baseline model is the property that agents face an additional source of informational heterogeneity: namely, uncertainty regarding how informed other agents might be about the nominal demand shock. Note then that the coefficient γ parameterizes the level of this heterogeneity: when $\gamma = 1$, all agents have the same precisions, and this fact is common knowledge; when instead $\gamma \in (1/2, 1)$, different agents have different precisions, and each agent is uncertain about the distribution of precisions in the rest of the population.

Note further that knowing the shock s_t would not help any agent improve his forecast of θ_t. This is simply because the mean and precision of this forecast depend merely on the agent's own

κ, not on the κ's of other agents. Nevertheless, the agent would love to know s_t because this could help him improve his forecast of the actions that other agents are taking in equilibrium. Indeed, since an agent's own expectation of θ_t depends on both his x_{it} and his κ_{it}, it is a safe guess that the equilibrium choice of the agent also depends on both x_{it} and κ_{it} and therefore that the aggregate price level depends both on θ_t and on κ_t. It then follows that agents face uncertainty about the aggregate price level, not only because they don't know the underlying nominal shock, but also because they don't know how much other agents are informed about that shock. Finally, note that, while the uncertainty about θ_t matters even when agents' actions are strategically independent ($\alpha = 0$), the latter type matters only when actions are interdependent ($\alpha \neq 0$). This highlights the distinctive nature of the additional source of uncertainty that we have introduced in this section.

To better appreciate this point, it is useful to study the stochastic properties of the hierarchy of beliefs about θ. Let E_t^1 denote the cross-sectional average of $\mathbb{E}_{it}[\theta_t]$ conditional on the current state for the precisions being s_t. Next, for any $k \geq 2$, let E_t^k denote the cross-sectional average of $\mathbb{E}_{it}[E_t^{k-1}]$; that's the k^{th}-order average beliefs. Clearly, all these average beliefs are functions of the realized aggregate state, (θ_t, s_t). Finally, let $\bar{E}_t^k$ denote the mean of the k^{th}-order belief if we fix the nominal shock and integrate over the precision shocks. It is easy to check that

$$\bar{E}_t^k = \eta_k \theta_t + (1 - \eta_k)\theta_{t-1},$$

for some constant η_k. The constant $\eta_k \equiv \partial \bar{E}_t^k / \partial \theta_t$ thus identifies the sensitivity of the k-th order belief to the underlying nominal shock. The response of the price level to the underlying nominal shock is determined by the sensitivities $\{\eta_k\}$.[9]

We illustrate the behavior of the hierarchy of beliefs in Figures 6 and 7. In Figure 6, we focus on the impact of the precision of information when this precision is common knowledge. We thus restrict $\kappa_h = \kappa_l = \kappa$ (in which case γ becomes irrelevant) and consider how the sensitivities of the beliefs to the shock vary with κ (which now identifies the common precision of information). We then observe the following qualitative properties. First, the hierarchy of beliefs about θ_t converges

[9] To see this more clearly, let for a moment $\lambda = 0$ (that is, abstract from adjustment frictions). The price set by each firm is then given by

$$p_{it} = \mathbb{E}_{it}[(1 - \alpha)\theta_t + \alpha p_t].$$

It follows that the price level satisfies

$$p_t = (1 - \alpha)\left(E_t^1 + \alpha E_t^2 + \alpha^2 E_t^3 + ...\right),$$

which implies that the response of the price level to nominal shocks is dictated by the response of the series E_t^k to θ_t.

Figure 6

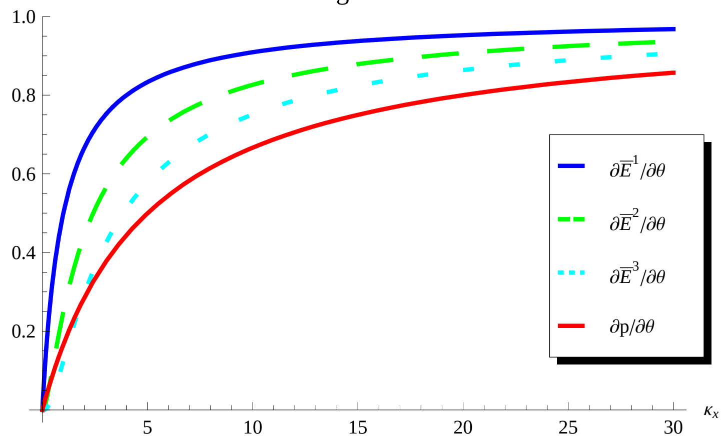

Figure 7

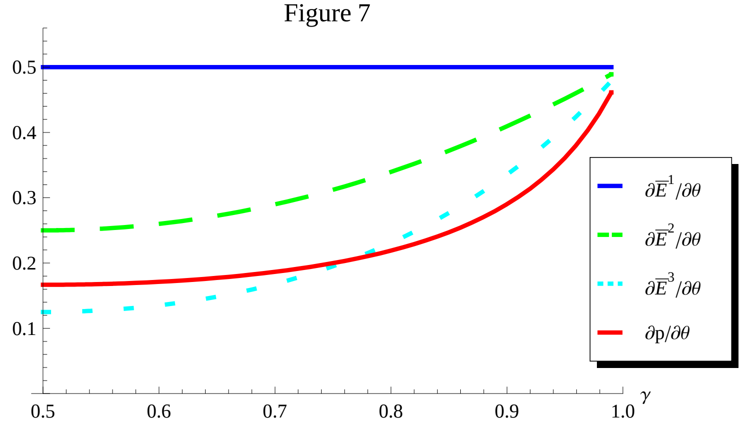

to the common prior expectation, θ_{t-1}, as the signals become uninformative: for all k, $\eta_k \to 0$ as $\kappa \to 0$. Second, the beliefs converge to the true underlying state, θ_t, as the signals become perfect: for all k, $\eta_k \to 1$ as $\kappa \to \infty$. Finally, whenever the signals are informative but not perfect, higher-order beliefs are more anchored towards the prior than lower-order beliefs: for any $\kappa \in (0, \infty)$, $1 > \eta_1 > \eta_2 > ... > 0$.

In Figure 7, we turn our focus to the impact of the uncertainty regarding the precision of others' information. In particular, we let $\kappa_h > \kappa_l$ and consider how the beliefs vary with the coefficient γ (which parameterizes the heterogeneity of information regarding the underlying precision state). We then observe that η_1 is invariant to γ, while η_2 and η_3 increase with γ. That is, the sensitivity of first-order beliefs to the nominal shock is independent of γ, while the sensitivities of higher-order beliefs increase with γ.

Along with the fact that the price level depends not only on first-order but also on higher-order beliefs, we can expect that γ should affect the response of the price level to the nominal shock even though it does not affect the response of first-order beliefs. Indeed, following similar steps as in the baseline model, we can solve for the equilibrium as follows.

Proposition 2. *(i) There exists an equilibrium in which the pricing strategy of a firm is given by*

$$p_{i,t} = \begin{cases} b_1 p_{t-1} + b_{2,h} x_{i,t} + b_{3,h} \theta_{t-1} & \text{if } \kappa_{i,t} = \kappa_h \\ b_1 p_{t-1} + b_{2,h} x_{i,t} + b_{3,h} \theta_{t-1} & \text{if } \kappa_{i,t} = \kappa_l \end{cases}$$

while the aggregate price level is given by

$$p_t = \begin{cases} c_1 p_{t-1} + c_{2,h} \theta_t + c_{3,h} \theta_{t-1} & \text{if } s_t = h \\ c_1 p_{t-1} + c_{2,h} \theta_t + c_{3,h} \theta_{t-1} & \text{if } s_t = l \end{cases}$$

for some coefficients $(b_1, b_{2,h}, b_{2,l}, b_{3,h}, b_{3,l})$ *and* $(c_1, c_{2,h}, c_{2,l}, c_{3,h}, c_{3,l})$.

(ii) Let $c_2 \equiv \frac{1}{2}(c_{2,h} + c_{2,l})$ *and* $c_3 \equiv \frac{1}{2}(c_{3,h} + c_{3,l})$ *be the mean sensitivity of the price level to the current and past nominal shock, averaging across the precision states. The equilibrium value of* c_1 *does not depend on* γ *and is identical to that in the baseline model, while the equilibrium values of* c_2 *and* c_3 *depend on* γ *if and only if* $\alpha \neq 0$.

This result is also illustrated in Figure 7, which plots the coefficient $c_2 \equiv \mathbb{E}[\partial p_t / \partial \theta_t]$ as a function of γ. We see that as γ decreases, the sensitivity of the first-order beliefs to the current nominal shock stays constant, while the sensitivity of the price level decreases. Ss anticipated, this is because a lower γ decreases the sensitivity of second- and higher-order beliefs.

To recap, the example of this section has highlighted how, even if one were to fix the sensitivity of the firms forecasts to the underlying nominal shock, one could still have significant freedom in how higher-order beliefs, and thereby equilibrium prices, respond to the shock. This is important for understanding the quantitative implications of incomplete information: to estimate the degree of price inertia caused by incomplete information, one may need direct or indirect information, not only about the firm's expectations about the underlying nominal shocks, but also about their higher-order expectations.

Finally, it is interesting to note how a variant of the model we have introduced here could generate the possibility that all firms are perfectly informed about the nominal shock, are free to adjust their prices fully, and yet find it optimal to adjust only partly. To see this, suppose that when the precision state is $s = h$ all firms get $\kappa_i = \infty$ (which means that their signals are perfectly informative); but when $s = l$, some firms get $\kappa_i = \infty$ and others get $\kappa_i = 0$ (which means that the signal is completely uninformative). Under this scenario, when the precision state is $s = h$, all firms are perfectly informed. However, this fact is not common knowledge. This is because each firm cannot tell whether she is perfectly informed because the state is $s = h$ or because the state is $s = l$ but was among the lucky ones to receive the perfectly precise signals. As a result, each firm must assign positive probability to the event that some other firms might not be informed and hence might not adjust their prices. But then because of strategic complementarity every firm will find it optimal not to adjust fully to the shock. It follows that there exist events where all firms are perfectly informed about the shock and nevertheless do not fully adjust their prices.

Of course, in this last example the possibility that firms are perfectly informed and yet do not respond perfectly to the shock can occur only with probability strictly less than one: the will also be events where some firms are relative uninformed and nevertheless find it optimal to respond quite a bit to the shock because they expect that other firms will be more informed. That is, this example cannot generate situations where in all events firms are perfectly informed and nevertheless expect other firms to be less informed. Indeed, based on the results of Kajii and Morris (1997a, 1997b) regarding the robustness of complete-information equilibria to the introduction of incomplete information, one can safely guess that if firms are nearly perfectly informed in almost all states of nature (formally, if the common prior assign probability near 1 to this event), then this fact will also be nearly common knowledge (formally, this event will be high common-p belief), which should then imply that all firms respond strongly to the shock in almost all states of nature. Nevertheless, the results of this section do highlight how quantifying the response of higher-order beliefs is essential for quantifying the response of prices to nominal shocks.

5 Heterogeneous priors

In this section we study how heterogeneous priors regarding the signals firms receive can affect the behavior of higher-order beliefs and thereby the response of prices to nominal shocks. In particular, we consider a system of heterogeneous priors that induce firms to behave in equilibrium as if they lived in a world where other firms were less informed about the underlying nominal shocks. This sustains a partially self-fulfilling equilibrium where firms react little to the underlying shock, even if they have nearly perfect information about it.

Apart from the introduction of heterogeneous priors, the setup is identical to our baseline Calvo model of Section 3. We again let θ_t follow an exogenous random walk process. In any given period, a firm may change its price with probability $1-\lambda$, in which case the price it chooses is a weighted average of all future target prices:

$$p_{i,t} = (1-\beta\lambda) \sum_{j=0}^{\infty} (\beta\lambda)^j \left[(1-\alpha) \mathbb{E}_{i,t}\theta_{t+j} + \alpha \mathbb{E}_{i,t} p_{t+j} \right] \tag{9}$$

As in the baseline model, each period firms learn perfectly the nominal demand of the previous period, θ_{t-1}, and receive a private signal of the current period's nominal demand:

$$x_{i,t} = \theta_t + \varepsilon_{i,t}.$$

However, firms disagree on the stochastic properties of the noise in their signals.

In particular, each firm believes that its own signal is an unbiased signal of θ_t. Specifically, firm i believes the error in its own private signal is drawn from the following distribution:

$$\varepsilon_{i,t} \sim \mathcal{N}(0, 1/\kappa_x).$$

At the same time, each firm believes that the private signals of all other firms are biased. Specifically, firm i believes that the errors in the private signals of all other firms are drawn independently from the following distribution:

$$\varepsilon_{j,t} \sim N\left(\delta_{i,t}, 1/\kappa_x\right) \;\; \forall j \neq i$$

where $\delta_{i,t}$ is the bias that firm i believes is present in the private signals of other firms. Finally, we assume that the perceived biases are negatively correlated with the innovation in the fundamental (the nominal shock). Specifically, we assume that, for all i and all t,

$$\delta_{i,t} = \delta_t \equiv -\chi\nu_t = -\chi(\theta_t - \theta_{t-1}),$$

where $\chi \in [0,1]$ is a parameter that controls the correlation of the bias with the innovation in the nominal shock. All these facts are common knowledge: the firms have heterogeneous priors, but these priors are common knowledge.

To understand the difference between the baseline model (which had assumed a common prior) and the current model (with allows for heterogeneous priors), it is useful to consider the beliefs of each firm about the *average* expectation of θ_t. In either model, each firm's own (first-order) expectation of the fundamental is

$$\mathbb{E}_{i,t}\theta_t = \frac{\kappa_x}{\kappa_x + \kappa_\theta} x_t + \frac{\kappa_\theta}{\kappa_x + \kappa_\theta} \theta_{t-1}$$

In the baseline model, this implied that each firm believed that the average first-order expectation in the rest of the population satisfied

$$\bar{E}_t^1 = \frac{\kappa_x}{\kappa_x + \kappa_\theta}\theta_t + \frac{\kappa_\theta}{\kappa_x + \kappa_\theta}\theta_{t-1}.$$

That is, the firm's second-order expectation was given by

$$\mathbb{E}_{i,t}\bar{E}_t^1 = \frac{\kappa_x}{\kappa_x + \kappa_\theta}\mathbb{E}_{i,t}\theta_t + \frac{\kappa_\theta}{\kappa_x + \kappa_\theta}\theta_{t-1}.$$

In contrast, now that firms have heterogeneous priors, each firm believes the average first-order expectation satisfies

$$\bar{E}_t^1 = \frac{(1-\chi)\,\kappa_x}{\kappa_x + \kappa_\theta}\theta_t + \frac{\kappa_\theta + \chi\kappa_x}{\kappa_x + \kappa_\theta}\theta_{t-1},$$

That is, the firm's second-order expectation is now given by

$$\mathbb{E}_{i,t}\bar{E}_t^1 = \frac{(1-\chi)\,\kappa_x}{\kappa_x + \kappa_\theta}\mathbb{E}_{i,t}\theta_t + \frac{\kappa_\theta + \chi\kappa_x}{\kappa_x + \kappa_\theta}\theta_{t-1},$$

Therefore, the heterogeneous priors that we have introduced in this section do not affect first-order beliefs, but they do affect second- and higher-order beliefs: the higher χ is, the more each firm believes that the beliefs of others will be less sensitive to innovations θ, even though its own belief is not affected.

We now examine how this affects equilibrium behavior. We conjecture once again an equilibrium in which the price set by a firm in period t is a linear function of $(p_{t-1}, \theta_{t-1}, x_t,)$:

$$p_{i,t} = P(p_{t-1}, \theta_{t-1}, x_{i,t}) = b_1 p_{t-1} + b_2 x_{i,t} + b_3 \theta_{t-1} \tag{10}$$

for some coefficients b_1, b_2, b_3. Accordingly, firm i expects that the price level will satisfy

$$p_t = \lambda p_{t-1} + (1-\lambda)\int P(p_{t-1}, \theta_{t-1}, x)dF_t(x) \tag{11}$$

where F_t is the cross-sectional distribution of signals as perceived by the typical firm. Our assumption regarding the heterogeneous priors implies that each firm thinks that the cross-sectional mean of the signals in the rest of the population is $(1-\chi)\theta_t + \chi\theta_{t-1}$. It follows that each firm expects the price level to satisfy

$$p_t = \lambda p_{t-1} + (1-\lambda)\,P(p_{t-1}, \theta_{t-1}, (1-\chi)\theta_t + \chi\theta_{t-1})$$

or equivalently

$$p_t = c_1 p_{t-1} + c_2\theta_t + c_3\theta_{t-1} \tag{12}$$

where

$$c_1 = \lambda + (1-\lambda)b_1, \quad c_2 = (1-\lambda)b_2(1-\chi) \quad c_3 = (1-\lambda)(b_3 + b_2\chi). \tag{13}$$

But now recall from the baseline model that, no matter what are the coefficients (c_1, c_2, c_3), the best response of a firm to (12) is to set a price as in (10), with the coefficients (b_1, b_2, b_3) defined by the solution to (8). We conclude that the equilibrium values of the coefficients (b_1, b_2, b_3) and (c_1, c_2, c_3) are now given by the joint solution of (8) and (13).

Note that χ enters only the conditions for c_2 and c_3 in (13), not the condition for c_1. It follows that the equilibrium value of c_1 (and hence also that of b_1) remains the same as in our baseline model (or, equivalently, as in the standard Calvo model). Moreover, the price process continues to be homogeneous of degree one, so that $c_1 + c_2 + c_3 = 1$. Finally, the equilibrium value of c_2 now satisfies

$$c_2 = \frac{\lambda(1-c_1)}{\lambda + c_1 \frac{\kappa_\theta + \chi\kappa_x}{\kappa_x(1-\chi)}}.$$

Comparing this with the corresponding condition for the baseline model, we observe that the equilibrium values of (c_1, c_2, c_3) for the present model coincide with those of the baseline model if the precision of information in that model is adjusted to the value $\tilde{\kappa}_x$ defined by

$$\frac{\tilde{\kappa}_x}{\kappa_\theta} \equiv \frac{(1-\chi)\,\kappa_x}{\kappa_x + \chi\kappa_\theta}, \tag{14}$$

which is clearly decreasing in χ.

This observation establishes a certain isomorphism between the present model and the baseline one: in the heterogeneous-prior economy we have introduced here, firms expect the price level to respond to the underlying nominal shock in the same way as in a variant of the baseline model in which the precision of information is decreased from κ_x to $\tilde{\kappa}_x$. Along with the fact that, because of strategic complementarity ($\alpha > 0$), the incentive of a firm to respond to its own information is lower the lower the expected response of the price level, we conclude that heterogeneous priors reduce the response of each firm to its own information about the underlying nominal shock.

Proposition 3. *In the equilibrium of the heterogeneous-priors economy, firms respond to their information in the same way as in an a common-prior economy in which the precision of information that other firms have is a decreasing function of χ. By implication, the sensitivity b_2 of a firm's price to its own signal about the underlying nominal shock is also decreasing in χ for given precision.*

To further appreciate this result, for the moment allow us to abstract from price rigidities ($\lambda = 0$) and consider the limit as $\kappa_x \to \infty$, meaning that each firm is (nearly) perfectly informed about the shock. In the common-prior world, this would have guaranteed that prices move one-to-one with the nominal shock and hence that the nominal shock has no real effect. But now let $\chi \to 1$ along with $\kappa_x \to \infty$ in such a way that the quantity $\tilde{\kappa}_x$ defined in (14) stays bounded away from ∞. In this limit, each firm is perfectly informed about the shock but expects other firms to respond as if they were imperfectly informed; firms therefore find it optimal to adjust their prices less than one-to-one in response to the nominal shock, thereby causing the shock to have a real effect, even though they are perfectly informed about the shock and there is no price rigidity.

The preceding analysis has focused on how heterogeneous priors affect the instantaneous response of prices to the underlying shock: in the model considered above, the dynamics of prices after the initial shock is driven solely by the Calvo mechanics, much alike the baseline model. However, heterogeneous priors can also affect these dynamics, to the extent that the perceived bias is persistent over time. To see this, suppose that bias δ_t follows an autoregressive process of the following form:

$$\delta_t = -\chi v_t + \rho \delta_t,$$

for some $\rho \in [0, 1)$. One can the easily extend the preceding analysis to show the following.

Proposition 4. *There exist coefficients* (b_1, b_2, b_3, b_4), *which depend on* (χ, ρ), *such that the equilibrium strategy of firm* i *is given by*

$$p_{i,t} = P(p_{t-1}, x_{it}, \theta_{t-1}, \delta_{t-1}) \equiv b_1 p_{t-1} + b_2 x_{it} + b_3 \theta_{t-1} + b_4 \delta_{t-1}$$

At this point, it is important to recognize that so far we have used the model to make predictions only about what the *firms* expect the price level to do and how they respond to their information about this shock—we have not used the model to make predictions about what *we*, as outside observers or "econometricians", expect the price level to do. This is where heterogeneous priors make things delicate. To analyze the equilibrium strategy of the firm as a function of their signals (their "types"), we do not need to take a stand on whether the firms signals are "truly" biased or not; we only need to postulate the system of their beliefs and then we can use the model to make predictions about how these beliefs will map into behavior. In contrast, to analyze the impulse response of the aggregate price level to the underlying shock, it no more suffices to characterize the mapping from beliefs to behavior; we also need to know on what is the mapping from the underlying nominal shock to the cross-sectional distribution of beliefs. In particular, we must now take a stand on how the cross-sectional average of signals relates to the underlying nominal shock.

Here, we could assume either that "econometrician" believes that the signals are biased, so that the cross-sectional average of x_{it} is $\theta_t + \delta_t$. or that she believes that they are unbiased, so that the cross-sectional average of x_{it} is θ_t. If we assume the former scenario, then the econometrician's law of motion of the price level will coincide with the one in the minds of the firms. In particular, it will be given by

$$p_t = (1-\lambda)p_{t-1} + \lambda P(p_{t-1}, \theta_t + \delta_t, \theta_{t-1}, \delta_{t-1}),$$

with the function P given by Proposition 4. If instead we assume the latter scenario, then the econometrician's law of motion for the price level will differ from the one in the mind of the firms. In particular, it will be given by

$$p_t = (1-\lambda)p_{t-1} + \lambda P(p_{t-1}, \theta_t, \theta_{t-1}, \delta_{t-1}),$$

with the function P given once again by Proposition 4. The only difference in these two law of motions in that the cross-sectional average of x_{it} is assumed to be $\theta_t + \delta_t$ in the first case and θ_t in the latter case.

For the remainder of the analysis, we will assume the latter scenario, keeping though in mind two properties. First, the two scenarios deliver similar qualitative properties as long as $\alpha > 0$. And second, the quantitative difference between the two scenarios vanishes as $\alpha \to 1$. Both of these properties are direct implications of strategic complementarity.

Assuming the second scenario and combining the stochastic process of θ_t and δ_t with the law of motion for the price level, we conclude that the dynamics of the economy, as seen from the perspective of the econometrician, are given by the following:

$$\begin{bmatrix} \theta_t \\ \delta_t \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \rho & 0 \\ \lambda b_2 + \lambda b_3 & \lambda b_4 & 1-\lambda+\lambda b_1 \end{bmatrix} \begin{bmatrix} \theta_{t-1} \\ \delta_{t-1} \\ p_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ -\chi \\ c_2 \end{bmatrix} v_t$$

with the coefficients (b_1, b_2, b_3, b_4) being determined by Proposition 4. We can then use this dynamic system, along with the fact that real output is $y_t = \theta_t - p_t$, to simulate the impulse responses of inflation and output to a positive innovation in v_t (of a size equal to the standard deviation of v_t).

Figures 8 and 9 illustrate these impulse responses for different values of χ and ρ. The baseline (common-prior) model corresponds to $\chi = \rho = 0$. As anticipated, we see that letting $\chi > 0$ but keeping $\rho = 0$ affects the impact effect of the innovation but not its persistence. In particular, in the period that the the innovation in nominal demand materializes, the response of inflation is dampened by letting $\chi > 0$, and by implication the positive effect on real output is amplified.

Figure 8: inflation

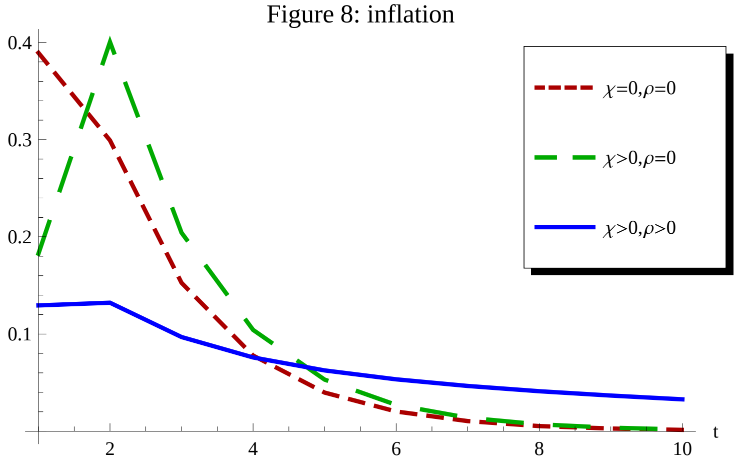

Figure 9: output

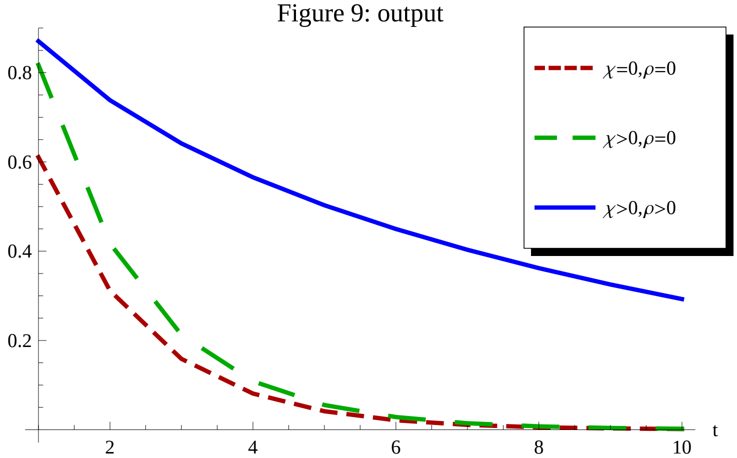

But as long as $\rho = 0$ the dynamics following this initial period are determined solely by the Calvo propagation mechanism and hence the persistence is the same as in the baseline model. This is because any discrepancy between either first-order or higher-order beliefs and the true value of the shock vanishes after the initial period. In contrast, letting $\rho > 0$ permits the discrepancy to persist in higher order beliefs even after it has vanished in first-order beliefs, thereby contributing to additional persistence in the real effects of the nominal shock.

To sum up, heterogeneous priors can help rationalize significant inertia the response of prices to changes in nominal demand simply by inducing inertia in the response of higher-order expectations. This is true no matter how high is the firms' precision of information, that is, the sensitivity of first-order beliefs to the nominal shock. Finally, note that the present model shares a bit of the flavor of the model with uncertain precisions that we considered in the previous section: there we focused on the possibility that firms may face uncertainty about the precision of other firms' information about the shock; here we showed how heterogeneous priors can induce firms to behave as if they expected other firms to have less precise information than themselves. In either case, the key is the behavior of higher-order beliefs as opposed to first-order beliefs.

6 Heterogenous priors and cost-push shocks

In the preceding two sections we highlighted how higher-order beliefs can induce inertia in the response of prices to nominal shocks in the economy. In so doing, we focused on the role of higher-order beliefs for the propagation of certain structural shocks. In this section we highlight how higher-order beliefs can themselves be the source of fluctuations in the economy.

For this purpose, we modify the variant of the previous section in the following way. Firms continue to have heterogeneous priors about their signals, but the perceived bias is no more correlated with the underlying nominal shock. Rather, the bias follows an independent stochastic process, given by

$$\delta_t = \rho \delta_{t-1} + \omega_t \tag{15}$$

where ω_t is a Normally distributed shock that is i.i.d. across time and independent of θ_τ for all τ.

Following similar steps as in the previous section, one can then show the following.

Proposition 5. *There exist coefficients* $(b_1, b_2, b_3, b_4, b_5)$ *such that the equilibrium strategy of firm* i *is given by*

$$p_{i,t} = b_1 p_{t-1} + b_2 x_{i,t} + b_3 \theta_{t-1} + b_4 \delta_{t-1} + b_5 \omega_t$$

Figure 10: inflation

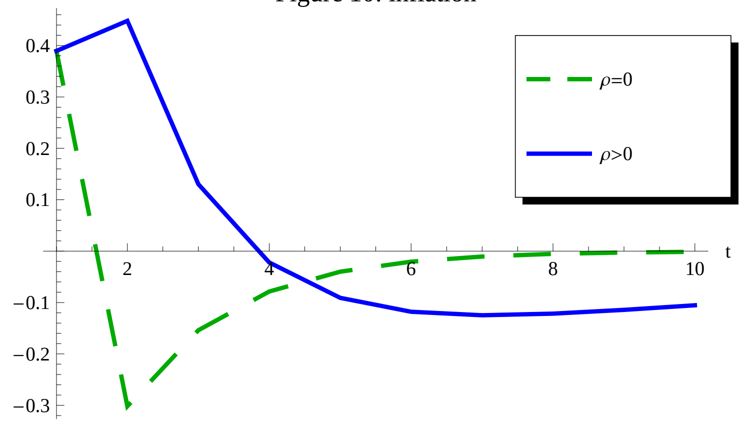

Figure 11: output

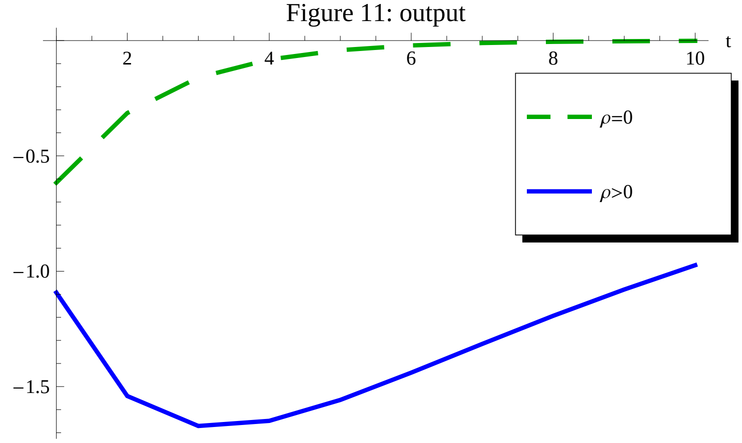

Taking once again the perspective of an econometrician who believes that the signals are unbiased, we obtain the impulse responses of the price level from the following dynamics:

$$\begin{bmatrix} \theta_t \\ \delta_t \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \rho & 0 \\ \lambda b_2 + \lambda b_3 & \lambda b_4 & 1-\lambda+\lambda b_1 \end{bmatrix} \begin{bmatrix} \theta_{t-1} \\ \delta_{t-1} \\ p_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \\ \lambda b_2 \end{bmatrix} v_t + \begin{bmatrix} 0 \\ 1 \\ \lambda b_5 \end{bmatrix} \omega_t$$

In Figures 10 and 11, we illustrate the impulse responses of inflation and real output to a positive innovation in ω_t. Such a shock causes firms to raise their prices even though aggregate nominal demand hasn't change. As a result, inflation increases and output contracts. The resulting fluctuations thus resembled to those often identified as the impact of "cost-push" shocks. When $\rho = 0$, this cost-push-like shock is transitory; the moderately persistent effects on real output are then due merely to the Calvo mechanics. When instead $\rho > 0$, the bias is itself persistent, which contributes to additional persistence in the real effects of the shock.

7 Conclusion

In this paper we studied how the combination of incomplete information and sticky prices may dampen the response of prices to nominal shocks. We did so through four models which progressively shifted focus from the precision of available information about the shocks to the behavior of higher-order beliefs. We thus sought to highlight that the quantifying the speed of learning—equivalently, the rate at which first-order adjust to the shock—does not suffice for quantifying the rate at which higher-order beliefs adjust, and therefore also does not suffice for quantifying the rate at which prices adjust.

We thus hope that more effort will be devoted to quantifying the behavior of higher-order expectations and their implications for the degree of price inertia at the macro level. In particular, we believe that a promising direction for future work is to explore the information that is contained in surveys of consumer and professional forecasts. Because these forecasts regard the key macroeconomic variables of interest for the agents, they may serve as sufficient statistics, or at least as good proxies, for the entire hierarchy of beliefs with regard to the behavior of these agents.

Appendix

Proof of Proposition 1. The condition that determines the equilibrium value of c_1 can be restated as

$$f(c_1, \lambda, \beta) = \alpha, \tag{16}$$

where

$$f(c_1, \lambda, \beta) \equiv \frac{(1 - \beta\lambda c_1)(c_1 - \lambda)}{c_1 (1 - \lambda)(1 - \beta\lambda)}.$$

As mentioned in the main text, there are two solutions to this equation—one with $c_1 > 1$ and $c_1 \in (\lambda, 1)$—and we focus on the non-explosive one.

Clearly, that the aforementioned equation is independent of κ_x and κ_θ, implying that c_1 is independent of the information structure. Moreover,

$$\frac{\partial c_1}{\partial \lambda} = -\frac{\partial f / \partial \lambda}{\partial f / \partial c_1} = \frac{c_1(1 - c_1)(1 - \beta c_1)(1 - \beta\lambda^2)}{\lambda(1 - \lambda)(1 - \beta\lambda)(1 - \beta c_1^2)} > 0$$

and

$$\frac{\partial c_1}{\partial \alpha} = \frac{1}{\partial f / \partial c_1} = \frac{\lambda(1 - \beta c_1^2)}{c_1^2(1 - \lambda)(1 - \beta\lambda)} > 0,$$

so that c_1 is increasing in both λ and α. Next, recall from the main text that c_2 satisfies

$$c_2 = \frac{\lambda(1 - c_1)}{\lambda + c_1 \frac{\kappa_\theta}{\kappa_x}}.$$

Since c_1 is independent of $(\kappa_x, \kappa_\theta)$, it is immediate that c_2 is increasing in κ_x/κ_θ; and since $c_3 = 1 - c_1 - c_2$, it is immediate that c_3 is decreasing in κ_x/κ_θ. Moreover, since the above expression for c_2 is independent of α for given c_1 and is decreasing in c_1, and since c_1 is itself increasing in α, it follows that c_2 is decreasing in α. Finally, the fact that c_2 is non-monotonic in λ and that c_3 is non-monotone in α can be establish by numerical example.

Proof of Proposition 2. The equilibrium can be characterized in a similar fashion as in the baseline model. First, by aggregating the strategy of the firms, we infer that the coefficients $(c_1, c_{2,h}, c_{2,l}, c_{3,h}, c_{3,l})$ must solve the following system:

$$\begin{aligned}
c_1 &= \lambda + (1 - \lambda) b_1 \\
c_{2,s} &= (1 - \lambda) \left[\gamma b_{2,s} + (1 - \gamma) b_{2,-s}\right] \\
c_{3,s} &= (1 - \lambda) \left[\gamma b_{3,s} + (1 - \gamma) b_{3,-s}\right]
\end{aligned}$$

where we use the convention that $-s = l$ when $s = h$ and $-s = h$ when $s = l$. Next, by taking the best response of the firm, we infer that the coefficients $(b_1, b_{2,h}, b_{2,l}, b_{3,h}, b_{3,l})$ solve the following system:

$$\begin{aligned}
b_1 &= ((1-\beta\lambda)\alpha + (\beta\lambda)b_1)\, c_1 \\
b_{2,s} &= [\ (1-\beta\lambda)\,(1-\alpha) + ((1-\beta\lambda)\alpha + (\beta\lambda)b_1)\,(\gamma c_{2,s} + (1-\gamma)c_{2,-s}) \\
&\quad +(\beta\lambda)\left(\tfrac{1}{2}(b_{2,l}+b_{2,h}) + (b_{3,l}+b_{3,h})\right)\ ]\,\frac{\kappa_x}{\kappa_x+\kappa_\theta} \\
b_{3,s} &= [\ (1-\beta\lambda)\,(1-\alpha) + ((1-\beta\lambda)\alpha + (\beta\lambda)b_1)\,(\gamma c_{2,s} + (1-\gamma)c_{2,-s}) \\
&\quad +(\beta\lambda)\left(\tfrac{1}{2}(b_{2,l}+b_{2,h}) + (b_{3,l}+b_{3,h})\right)\ ]\,\frac{\kappa_\theta}{\kappa_x+\kappa_\theta} \\
&\quad + ((1-\beta\lambda)\alpha + (\beta\lambda)b_1)\,(\gamma c_{3,s} + (1-\gamma)c_{3,-s})\,.
\end{aligned}$$

Clearly, c_1 and b_1 continue to be determined by the same equations as in the baseline model. Once again, we ignore the solution with $c_1 > 1$ and focus on the solution with $c_1 \in (0,1)$. Given this solution, the remainder of the conditions consist a linear system, which admits a unique solution for the coefficients $(b_{2,s}, b_{3,s}, c_{2,s}, c_{3,s})_{s\in\{h,l\}}$.

Proof of Proposition 3. In the main text we showed that the equilibrium values of (c_1, c_2, c_3) and (b_1, b_2, b_3) are determined by the solution to (8) and (13); that c_1 and b_1 continue to be determined as in the baseline model and are thus independent of χ; and that c_2 satisfies

$$c_2 = \frac{\lambda(1-c_1)}{\lambda + c_1\frac{\kappa_\theta+\chi\kappa_x}{\kappa_x(1-\chi)}},$$

which is decreasing in χ. Along with the fact that $c_2 = (1-\lambda)b_2(1-\chi)$, we get that

$$b_2 = \frac{\lambda(1-c_1)}{(1-\lambda)\left[\lambda + c_1\frac{\kappa_\theta}{\kappa_x} + (c_1-\lambda)\chi\right]},$$

which is also decreasing in χ, since $c_1 \in (\lambda, 1)$.

Proof of Propositions 4 and 5. To nest both the model of Sections 5 and 6, we let the bias be given by

$$\delta_t = -\chi v_t + \omega_t + \rho\delta_{t-1}.$$

We then conjecture an equilibrium in which the price set by a firm in period t is a linear function of $(p_{t-1}, x_t, \theta_{t-1}\delta_{t-1}, \omega_t)$:

$$p_{i,t} = P(p_{t-1}, x_{i,t}, \theta_{t-1}, \delta_{t-1}, \omega_t) = b_1 p_{t-1} + b_2 x_{i,t} + b_3\theta_{t-1} + b_4\delta_{t-1} + b_5\omega_t \tag{17}$$

for some coefficients b_1, b_2, b_3, b_4, b_5.

Our specification of the heterogeneous priors implies that each firm thinks that the cross-sectional mean of the signals in the rest of the population is $\bar{x}_t = \theta_t + \delta_t = (1-\chi)\theta_t + \chi\theta_{t-1} + \rho\delta_{t-1} + \omega_t$. It follows that each firm expects the price level to satisfy

$$p_t = \lambda p_{t-1} + (1-\lambda)\, P(p_{t-1}, \bar{x}_t, \theta_{t-1}, \delta_{t-1}, \omega_t)$$

or equivalently

$$p_t = c_1 p_{t-1} + c_2\theta_t + c_3\theta_{t-1} + c_4\delta_{t-1} + c_5\omega_t \tag{18}$$

where

$$\begin{aligned} c_1 &= \lambda + (1-\lambda)\, b_1, \quad c_2 = (1-\lambda)\, b_2\, (1-\chi)\,, \quad c_3 = (1-\lambda)\,(b_2\chi + b_3) \\ c_4 &= (1-\lambda)\,(b_2\rho + b_4)\,, \quad c_5 = (1-\lambda)\,(b_2 + b_5) \end{aligned} \tag{19}$$

Next, note that we may write the firm i's best response (9) as

$$p_{i,t} = (1-\beta\lambda)\left[(1-\alpha)\,\mathbb{E}_{i,t}\theta_t + \alpha\mathbb{E}_{i,t}p_t\right] + (\beta\lambda)\,\mathbb{E}_{i,t}p_{i,t+1}.$$

Given that the firm expects the price level to evolve according to (18), for the firm's best response to be consistent with our conjecture (17), it must be that the coefficients $(b_1, b_2, b_3, b_4, b_5)$ solve the following system:

$$\begin{aligned} b_1 &= \left[(1-\beta\lambda)\,\alpha + (\beta\lambda)\, b_1\right] c_1 \\ b_2 &= \left[(1-\beta\lambda)\,(1-\alpha+\alpha c_2) + (\beta\lambda) + (\beta\lambda)\,(b_1c_2 + b_2 + b_3 - \chi b_4)\right] \frac{\kappa_x}{\kappa_x + \kappa_\theta} \\ b_3 &= \left[(1-\beta\lambda)\,(1-\alpha+\alpha c_2) + (\beta\lambda) + (\beta\lambda)\,(b_1c_2 + b_2 + b_3 - \chi b_4)\right] \frac{\kappa_\theta}{\kappa_x + \kappa_\theta} \\ &\quad + (1-\beta\lambda)\,\alpha c_3 + (\beta\lambda)\, b_1 c_3 + (\beta\lambda)\, b_4\chi \\ b_4 &= (1-\beta\lambda)\,\alpha c_4 + (\beta\lambda)\, b_1 c_4 + (\beta\lambda)\, b_4\rho \\ b_5 &= (1-\beta\lambda)\,\alpha c_5 + (\beta\lambda)\, b_1 c_5 + (\beta\lambda)\, b_4 \end{aligned} \tag{20}$$

Combining conditions (19) and (20) gives us a system of equations which characterizes the equilibrium values for $(b_1, b_2, b_3, b_4, b_5, c_1, c_2, c_3, c_4, c_5)$.

It is immediate that the conditions that determine c_1 and b_1 are identical to those in the baseline model. As in the baseline model, we ignore the solution that has $c_1 > 1$ and focus on the solution that has $c_1 \in (0,1)$. Given this solution, the remaining conditions define a linear system, which has a unique solution for the remaining coefficients.

References

[1] Amato, Jeffrey, and Hyun Song Shin (2006), "Imperfect Common Knowledge and the Information Value of Prices," *Economic Theory* 27, 213-241.

[2] Angeletos, George-Marios, and Jennifer La'O (2008a), "Dispersed Information over the Business Cycle: Optimal Fiscal and Monetary Policy," work in progress.

[3] Angeletos, George-Marios, and Jennifer La'O (2008b), "Higher-Order Beliefs, Market Sentiment, and the Business Cycle," work in progress.

[4] Angeletos, George-Marios, and Alessandro Pavan (2007), "Efficient Use of Information and Social Value of Information," *Econometrica* 75:4, 1103-1142.

[5] Bacchetta, Philippe, and Eric van Wincoop (2005), "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?," *American Economic Review* 95.

[6] Barro, Robert (1976), "Rational Expectations and the Role of Monetary Policy," *Journal of Monetary Economics* 2, 1-32.

[7] Bils, Mark, and Pete Klenow (2004), "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy* 112, 947-985.

[8] Calvo, Guillermo (1983), "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics* 12, 383-398.

[9] Collard, Fabrice, and Harris Dellas (2005), "Misperceived vs Unanticipated Money: A Synthesis," CNRS-GREMAQ/University of Bern mimeo.

[10] Hellwig, Christian (2005), "Heterogeneous Information and the Welfare Effects of Public Information Disclosures," UCLA mimeo.

[11] Kajii, Atsushi, and Stephen Morris (1997a), "The Robustness of Equilibria to Incomplete Information," *Econometrica* 65, 1283-1309.

[12] Kajii, Atsushi, and Stephen Morris (1997a), "Common p-Belief: The General Case," *Games and Economic Behavior* 18, 73-82.

[13] Klenow, Pete, and Jonathan Willis (2007), "Sticky Information and Sticky Prices," *Journal of Monetary Economics* 54, 79-99.

[14] Lorenzoni, Guido (2008), "A Theory of Demand Shocks," MIT mimeo.

[15] Lucas, Robert E., Jr. (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4, 103-124.

[16] Mackowiak, Bartosz, and Mirko Wiederholt (2007) "Optimal Sticky Prices under Rational Inattention," Humboldt University/Northwestern University mimeo.

[17] 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, 1295-1328.

[18] Morris, Stephen, and Hyun Song Shin (2002), "The Social Value of Public Information," *American Economic Review* 92, 1521-1534.

[19] Morris, Stephen, and Hyun Song Shin (2006), "The Inertia of Forward-Looking Expectations," *American Economic Review* 96, 1521-1534.

[20] Phelps, Edmund (1970), "Introduction: The New Microeconomics in Employment and Inflation Theory," in A.A. Alchian, C.C. Holt et al., *Microeconomic Foundations of Employment and Inflation Theory,* New York: Norton.

[21] Reis, Ricardo (2006), "Inattentive Producers," *Review of Economic Studies* 73, 793-821.

[22] Scheinkman, Jose, and Wei Xiong (2003), "Overconfidence and Speculative Bubbles," *Journal of Political Economy.*

[23] Sims, Christopher (2003), "Implications of Rational Inattention," *Journal of Monetary Economics* 50, 665-690.

[24] Townsend, Robert M. (1983), "Forecasting the Forecasts of Others," *Journal of Political Economy* 91, 546-588.

[25] Woodford, Michael (2003), "Imperfect Common Knowledge and the Effects of Monetary Policy," in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton University Press.