

NBER WORKING PAPER SERIES

INCOMPLETE INFORMATION, HIGHER-ORDER BELIEFS AND PRICE INERTIA

George-Marios Angeletos
Jennifer La'O

Working Paper 15003
<http://www.nber.org/papers/w15003>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2009

This paper was prepared for the 2008 SNB/JME conference on “Incomplete Information and Monetary Policy,” hosted by the Gerzensee Study Center, Switzerland. We received useful comments from multiple conference participants. We are particularly grateful to our discussants, Stephen Morris and Elmar Mertens, and to the organizers of the conference, Sergio Rebelo and Robert King. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by George-Marios Angeletos and Jennifer La'O. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Incomplete Information, Higher-Order Beliefs and Price Inertia
George-Marios Angeletos and Jennifer La'O
NBER Working Paper No. 15003
May 2009
JEL No. D8,E1,E3

ABSTRACT

This paper investigates how incomplete information impacts the response of prices to nominal shocks. Our baseline model is a variant of the Calvo model in which firms observe the underlying nominal shocks with noise. In this model, the response of prices is pinned down by three parameters: the precision of available information about the nominal shock; 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 highlight that this synthesis provides only a partial view of the role of incomplete information. In general, the precision of information does not pin down the response of higher-order beliefs. Therefore, one cannot quantify the degree of price inertia without additional information about the dynamics of higher-order beliefs, or the agents' forecasts of inflation. We highlight the distinct role of higher-order beliefs with three extensions of our baseline model, all of which break the tight connection between the precision of information and higher-order beliefs featured in previous work.

George-Marios Angeletos
Department of Economics
MIT E52-251
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER
angelet@mit.edu

Jennifer La'O
MIT Department of Economics
jenlao@mit.edu

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 the sources and the propagation of the business cycle, as well as 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 convenient time-dependent models 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 imperfect information about these shocks. The older literature formalized this imperfection as a geographical dispersion of the available information (Lucas, 1972, Barro, 1976); more recent contributions have proposed infrequent updating of information (Mankiw and Reis, 2002; Reis, 2006) or rational inattention (Sims, 2003; Woodford, 2003, 2008; Machowiak and Wiederholt, 2008). One way or another, though, the key driving force behind price inertia is that firms happen, or choose, to have imperfect knowledge of the underlying nominal shocks.

The starting point of this paper is a bridge between these two approaches. In particular, our baseline model is a hybrid of Calvo (1983), Morris and Shin (2002), 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, similarly to Morris-Shin and Woodford.

Within this baseline 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 the underlying nominal shocks (equivalently, the level of noise in the firms’ signals of these shocks); the frequency of price adjustment; and the degree of strategic complementarity in pricing decisions. That all three parameters should matter is obvious; but their interaction is also interesting. 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 with regard to how frictions in either price adjustment or information about the underlying nominal shocks impact the response of prices to these shocks. This synthesis has its own value, as it provides a simple and tractable incomplete-information version of the Calvo model that could readily be taken to the data. Nevertheless, this synthesis is not the main contribution of the paper. Rather, the main contribution of the paper is, first, to highlight that the aforementioned lessons miss the distinct role that higher-order beliefs play in the dynamics of price adjustment and, second, to show how this distinct role can be parsimoniously accommodated within our Calvo-like framework.

The basic idea behind our contribution is simple. The precision of the firms' information about the underlying nominal shock identifies how fast the firms' forecasts of the shock respond to the true shock: the more precise their information, the faster their forecasts converge to the truth. However, this need not also identify how fast the forecasts of the forecasts of others may adjust. 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.

To better understand this point, 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:

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

where \bar{E}^k denotes the k^{th} -order average forecast of θ .¹ It then follows that the response of price level p to an innovation in θ depends on the response of the entire sequence of different orders of beliefs, $\{\bar{E}^k\}_{k=1}^{\infty}$, to that shock.

¹The k^{th} -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.

When information is perfect or at least commonly shared, then $\bar{E}^k = \bar{E}^1$ for all k . It then follows that the response of prices to an innovation in θ depends merely on the response of first-order beliefs, which in turn is pinned down by the precision of the available information about θ . When, instead, information is dispersed, higher-order beliefs need not coincide with first-order beliefs.

The potential role of higher-order beliefs has been noted before by Morris and Shin (2002), Woodford (2003), and others. However, in the standard Gaussian example used in the pertinent 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 information about the underlying nominal shock remains the key determinant of the response of the price level to the shock. However, once one goes away from the standard Gaussian example, this tight connection between first- and higher-order beliefs can break—and the break can be quite significant.

We highlight the crucial and distinct role of higher-order beliefs with three variants of our baseline model. All three variants retain the combination of infrequent price adjustment (as in Calvo) and noisy information about the underlying shocks (as in Morris-Shin and Woodford), but differentiate in the specification of higher-order beliefs.

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 necessarily 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 less precise information than what they themselves have. This in turn helps rationalize why equilibrium prices may adjust very slowly to the underlying nominal shocks even if the frequency of price adjustment is arbitrarily high and each firm has arbitrarily precise information about the underlying nominal shock. Once again, the key is the inertia of higher-order beliefs; heterogeneous priors is a convenient modeling device.

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 its 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 firms’ information (beliefs) about these shocks, but also about the stochastic properties of their forecasts of the forecasts of others (higher-order beliefs).

Because the hierarchy of beliefs is an infinitely dimensional object, incorporating the distinct role of higher-order beliefs in macroeconomic models may appear to be a challenging task. Part of the contribution of the paper is to show that this is not the case. All the models we present here are highly parsimonious and nevertheless allow for rich dynamics in higher-order beliefs.

The rest of the paper is organized as follows. Section 2 discusses the relation of our paper to the literature. Section 3 studies our baseline model, which introduces incomplete information in the Calvo model. Section 4 studies the variant with uncertainty about the precisions of one another. Section 5 studies the variant that with heterogeneous priors. Section 6 turns to cost-push shocks. Section 7 concludes with suggestions for future research. All proofs are in the Appendix.

2 Related literature

The macroeconomics literature on informational frictions has a long history, going back to Phelps (1970), Lucas (1972, 1975), Barro (1976), King (1983) and Townsend (1983). Recently, this literature has been revived by Mankiw and Reis (2002), Morris and Shin (2002), Sims (2003), Woodford (2003), and subsequent work.² This paper contributes to this literature in two ways: first, by studying the interaction of incomplete information with price rigidities within the Calvo model; second, and most importantly, by furthering our understanding of the distinct role of higher-order

²See, e.g., Amato and Shin (2006), Angeletos and La’O (2008, 2009a, 2009b), Angeletos and Pavan (2007), Bacchetta and Wincoop (2005), Collard and Dellas (2005), Hellwig (2005), Lorenzoni (2008, 2009), Mackowiak and Wiederholt (2008, 2009), Nimark (2007, 2008), and Reis (2006).

beliefs. Closely related in this respect is Angeletos and La'O (2009a), which emphasizes how dispersed information has very distinct implications for the business cycle than uncertainty about the fundamentals.

Our paper is highly complementary to the papers by Woodford (2003) and Morris and Shin (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.³ However, by adopting the convenience of a popular but very specific Gaussian 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 nominal 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 firms are nearly perfectly informed about these shocks (in which case first-order beliefs respond nearly one-to-one with the shock).

Closely related are also the papers by Nimark (2008) and Klenow and Willis (2007). The former paper studies a similar framework as ours, namely a Calvo model with incomplete information, along with a more complex learning dynamics: unlike our paper and rather as in Woodford (2003), the underlying shocks do not become common knowledge after one-period delay. The latter paper studies a menu-cost model with sticky information. Much alike our baseline model, both papers study the interaction of price rigidities and informational frictions. However, these papers do disentangle the role of higher-order beliefs as we do in this paper. In particular, the quantitative importance of incomplete information in these papers is tied to the precision of information about the underlying shocks at the time of price changes. In contrast, our paper shows how one can disentangle the dynamics of higher-order beliefs from the speed of learning, and uses this to argue that significant price inertia at the macro level can be consistent with both significant price flexibility at the micro level and fast learning about the underlying nominal shocks.

³A similar role of higher-order beliefs has been highlighted by Allen, Morris and Shin (2005), Angeletos, Lorenzoni and Pavan (2007) and Bacchetta and Wincoop (2005) within the context of financial markets.

Finally, the heterogeneous-priors variant of this paper builds on Angeletos and La’O (2009b). In that paper we consider a real-business-cycle model in which firms have dispersed information about the underlying productivity shocks. We then show how dispersed information opens the door to a certain type of sunspot-like fluctuations—i.e., fluctuations that cannot be explained by variation in *either* the underlying economic fundamentals or the firms’ beliefs about these fundamentals. These fluctuations obtain also under a common prior, but are easier to model with heterogeneous priors. The present paper complements this other work by illustrating how these type of fluctuations can take the form of cost-push shocks in a new-keynesian model, and how heterogeneous priors can also help rationalize significant inertia in the response of prices to nominal shocks.

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 $t \in \{0, 1, 2, \dots\}$. 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 preferences of the household are given by $\sum_t \beta^t U(C_t, N_t)$, with

$$U(C_t, N_t) = \log C_t - N_t,$$

where $\beta \in (0, 1)$ is the discount rate, N_t is the labor supplied by the household,

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

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.

By the resource constraint for each commodity, we have that $Y_{i,t} = C_{i,t}$ for all i and therefore aggregate output is given by $Y_t = C_t$. Finally, aggregate nominal demand is given by the following quantity-theory or cash-in-advance constraint:

$$P_t C_t = \Theta_t.$$

Here, Θ_t denotes the level of aggregate nominal demand (aggregate nominal GDP), is assumed to be exogenous, and defines the “monetary shock” of our model.

In what follows, we use lower-case variables for the logarithms of the corresponding upper-case variables: $\theta_t \equiv \log \Theta_t$, $y_t \equiv \log Y_t$, $p_t \equiv \log P_t$, 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. Let $\mathcal{I}_{i,t}^1$ and $\mathcal{I}_{i,t}^2$ denote the information set of firm i during, respectively, the morning and the evening of period t . Information about the current level of nominal demand is imperfect during the morning but perfect during the evening. 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 of firm i in the morning of period t is therefore given by $\mathcal{I}_{i,t}^1 = \mathcal{I}_{i,t-1}^2 \cup \{x_{i,t}, a_{i,t}\}$, while her information in the evening of the same period is given by $\mathcal{I}_{i,t}^2 = \mathcal{I}_{i,t-1}^1 \cup \{\theta_t\}$. We will see in a moment that, in equilibrium, p_{t-1} and y_{t-1} are functions of $\{\theta_{t-1}, \theta_{t-2}, \dots\}$; 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 difference if we had included the realized level of a firm's demand into its evening information set.

The assumption that θ_t becomes common knowledge at the end of each period is neither random nor inconsequential. If we wished information about θ_t to remain dispersed after the end of period t , we would need somehow to limit the aggregation of information that takes place through commodity markets. That would require a more decentralized trading structure and would complicate the necessary micro-foundations. Furthermore, the dynamics would now become much less tractable: as in Townsend (1983), Nimark (2008), and others, firms would now have to keep track of the entire history of their information in order to forecast the forecasts of others, and the equilibrium dynamics would cease to have any simple recursive structure. Here, we avoid all these complications, and keep the analysis highly tractable, only by assuming, as in Lucas (1972), that θ_t becomes common knowledge after a short delay. However, it is important to recognize that, in so doing, we impose a fast convergence of beliefs about the past shocks and also rule out any heterogeneity in the agents' expectations of future shocks beyond the one in their beliefs about the current shock. One may expect that relaxing these properties would add to even more inertia, both because firms would learn more slowly (Woodford, 2003) and because expectations of future shocks could be more anchored to public information (Morris and Shin, 2006).

Price-setting behavior. Consider for a moment the case where prices are flexible and the current nominal shock is common knowledge at the moment firms set prices. The optimal price set by firm i in period t is then 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

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

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 (i.e., the flexible-price full-information optimal 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

$$\alpha \equiv 1 - \frac{1}{\epsilon + (1 - \epsilon)\eta} \in (0, 1)$$

defines the degree of strategic complementarity in pricing decisions, and where

$$\xi_{i,t} \equiv \frac{\epsilon}{\epsilon + (1 - \epsilon)\eta} \mu - \frac{1}{\epsilon + (1 - \epsilon)\eta} a_{i,t}$$

is simply a linear transformation of the idiosyncratic productivity shock. We henceforth 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 common knowledge in the begging of period t , the firm would set $p_{i,t} = p_{i,t}^*$ in all periods and states. However, we have assumed that firms have only imperfect information. Moreover, 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 its current expectation of a weighted average of the current and all future target prices:

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

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 .⁴

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

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

In the remainder of the paper, we treat condition (3) as if it were an exogenous behavioral rule, with the understanding though that this rule is actually fully microfounded 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 Perfect-Bayesian Equilibrium.⁵ 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 $(p_{t-1}, \theta_{t-1}, x_{i,t}, \xi_{i,t})$:

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

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

⁴To be precise, condition (2) should have been written as $p_{i,t} = \text{const} + \mathbb{E}_{i,t} \left[(1 - \beta\lambda) \sum_{j=0}^{\infty} (\beta\lambda)^j p_{t+j}^* \right]$, where const is an endogenous quantity that involves second-order moments and that emerges due to risk aversion. However, under the log-normal structure of shocks and signals that we have assumed, these second-order moments are invariant with either the shock or the information of the firms, and it is thus without any loss of generality to ignore the aforementioned constant.

⁵We can safely ignore out-of-equilibrium paths by assuming that firms observe (at most) the cross-section distribution of prices, which guarantees that no individual deviation is detectable.

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

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

where

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

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

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

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

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

Next, note that

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

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:

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

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 part (ii) 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 widely available.⁶ We accordingly set $\beta = .95$ (which




⁶A qualification is due here. The fact that these data are widely available suggests that most agents are likely to be well informed about them. This in turn implies that their first-order beliefs are likely to converge to the truth very fast. However, to the extent that this fact is not common knowledge, it is possible that higher-order beliefs do not converge as fast, which could contribute to further inertia in the response of prices.

means a discount rate of about 1% per quarter), $\lambda = .20$ (which means a probability of price change equal to 1/3 per quarter), and $\alpha = .85$ (which means a quite strong complementarity in pricing decisions); these values are consistent with standard calibrations of the Calvo model. Lacking any obvious estimate of the precision of information about the underlying shocks, we set $\kappa_x/\kappa_\theta = 1$; this means that the variance of the forecast error of the typical firm about the current innovation in nominal demand is one half the variance of the innovation itself.⁷



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 c_1 is increasing in λ , c_3 is decreasing in λ , and c_2 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 baseline model is the same as in the standard Calvo model. This property, however, hinges on our assumption that the nominal shock becomes common knowledge only with a delay of one period. If we increase the length of this delay, then we can obtain more persistence, similarly to Woodford (2003) or Nimark (2008), but only till the shock becomes common knowledge; after that point, any subsequent persistence is driven solely by the Calvo rigidity. The property, then, that c_1 is increasing in the Calvo parameter λ should 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).

⁷Note that only the ratio κ_x/κ_θ , and not the absolute values of κ_x and κ_θ , matter for the equilibrium coefficients c_1 , c_2 , and c_3 . In other words, once we fix κ_x/κ_θ , κ_θ is only a scaling parameter. By implication, the impulse responses of the economy to a one-standard-deviation change in v are invariant to κ_θ .



Where the incompleteness of information has a bite in our baseline model is on the coefficients c_2 and c_3 , which characterize, respectively, the price impact of the current and the past shock 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 firm'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 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 $\beta = .95$, $\lambda = .20$ 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 perfect information (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 κ_x/κ_θ for all subsequent periods.

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 firms initially have little information about the innovation and can rise in the early phases of learning, but once firms 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 firms 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 firms to face uncertainty regarding the precision of information that other firms may have regarding nominal demand. In particular, we introduce a second aggregate state variable, which permits us to capture uncertainty about the average precision of available information in the cross-section of the economy.

This new state variable is modeled as a binary random variable, $s_t \in \{h, l\}$, which is i.i.d. over time and independent of the nominal shock θ_t , and which takes 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 firms, with $1/2 < \gamma < 1$ and $0 < \kappa_l < \kappa_h$. The "type" of a firm is now given by the pair $(x_{i,t}, \kappa_{i,t})$, where $x_{i,t} = \theta_t + \varepsilon_{i,t}$ is the particular signal the firm receives about the current nominal shock, $\varepsilon_{i,t} \sim \mathcal{N}(0, 1/\kappa_{i,t})$ is the noise in this signal, and $\kappa_{i,t}$ is its precision. The latter is specific to the firm and is contingent on the new state variable s_t as follows: when this state is $s_t = h$, the precision of firm i is $\kappa_{i,t} = \kappa_h$ with probability γ and $\kappa_{i,t} = \kappa_l$ with probability $1 - \gamma$; and, symmetrically, when the state is $s_t = l$, the precision of firm i is $\kappa_{i,t} = \kappa_l$ with probability γ and $\kappa_{i,t} = \kappa_h$ with probability $1 - \gamma$. Finally, because these realizations are independent across the firms, γ is also the fraction of the population whose signals have precision κ_s when the state is s , for $s \in \{h, l\}$.

Note that a firm knows his own $\kappa_{i,t}$, but not the underlying state s_t . A firm's $\kappa_{i,t}$ thus serves a double role: it is both the precision of the firm's own information about the nominal shock θ_t 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 firms face an additional source of

informational heterogeneity: they face uncertainty regarding how informed other firms might be about the nominal demand shock. Note then that the coefficient γ parameterizes the level of this heterogeneity: when $\gamma = 1$, all firms have the same precisions, and this fact is common knowledge; when instead $\gamma \in (1/2, 1)$, different firms have different precisions, and each firm is uncertain about the distribution of precisions in the rest of the population.


Furthermore, note that knowing the state variable s_t would not help any firm improve his forecast of the nominal shock θ_t . This is simply because belief of a firm about the nominal shock depends only on its own precision (which the firm knows), not on the precisions of other firms (which the firm does not know). Nevertheless, the firm would love to know s_t because this could help him improve his forecast of the forecasts and actions of other firms in equilibrium. Indeed, since a firm'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 firm 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 firms face uncertainty about the aggregate price level, not only because they do not know the underlying innovation in θ_t , but also because they don't know how precisely other firms are informed about this shock. Finally, note that, while the uncertainty about θ_t matters for individual pricing behavior, and hence for aggregate prices, even when firms' pricing decisions are strategically independent ($\alpha = 0$), the uncertainty about s_t matters only when their pricing decisions 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 current and past nominal shocks and the current precision state s_t . Finally, let \bar{E}_t^k denote the expectation of the k^{th} -order average belief conditional on the nominal shocks alone (that is, averaging across the two possible s_t states). 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 average belief to the underlying nominal shock. As anticipated in the Introduction, the response of the price level to the underlying nominal shock is determined by the sensitivities $\{\eta_k\}$.⁸

⁸The discussion in the Introduction had abstracted from price rigidities (i.e., it had imposed $\lambda = 0$), but the



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

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.⁹

insight is clearly more general.

⁹As mentioned earlier, in both our baseline model and in all three variants of it, we identify the equilibrium as the fixed point of the best-response condition (3). For the current variant, this is with some abuse, since the non-Gaussian nature of the information structure implies that this condition is not exact: it is only a log-linear approximation. However, the properties that first-order beliefs do not depend on γ , while higher-order beliefs and hence equilibrium

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. As 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 outcomes do depend on γ , do not hinge on this approximation. Finally, keep in mind that condition (3) is exact in either the baseline model or the two other variants that we consider subsequently.

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 she 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: there will also be events where some firms are relatively 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 (1997) regarding the robustness of complete-information equilibria to the introduction of incomplete information, one can safely guess that if the common prior assigns probability near 1 to the set of events where the firms are nearly perfectly informed about the underlying nominal shock, then with probability near 1 the response of prices to the underlying nominal shock will be nearly the same as in the common-knowledge benchmark. 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 the underlying nominal shocks. In particular, we consider a system of heterogeneous priors that induces 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:

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

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:

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

where $\delta_{i,t}$ is the bias that firm i believes to be 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 perceived bias with the innovation in the nominal shock. Finally, these perceptions are commonly understood and mutually accepted: the firms have agreed to disagree.¹⁰

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

¹⁰We have used a related heterogeneous-priors specification in Angeletos and La'O (2009b), albeit within a different context and for different purposes. We refer the reader to that paper for a more thorough discussion on the modeling role of heterogeneous priors: they are convenient, but need not be strictly necessary for the type of effects we document. For example, in that paper we document how a certain type of sunspot-like fluctuations can obtain with either heterogeneous priors or a common prior, but the former maintain a higher level of tractability.

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 $(p_{t-1}, \theta_{t-1}, x_t)$:

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

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

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

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

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

where

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

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 last condition 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

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

which is clearly decreasing in χ . This observation in turn establishes a certain isomorphism between the present model and the baseline one: in the heterogeneous-prior economy that we have introduced here, firms expect the price level to respond to the underlying nominal shock in the same way as in a common-prior economy that is identical to the one in our baseline model except for the fact that the precision of available 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 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 $\kappa_x \rightarrow \infty$, 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 $\chi \rightarrow 1$ along with $\kappa_x \rightarrow \infty$ 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 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

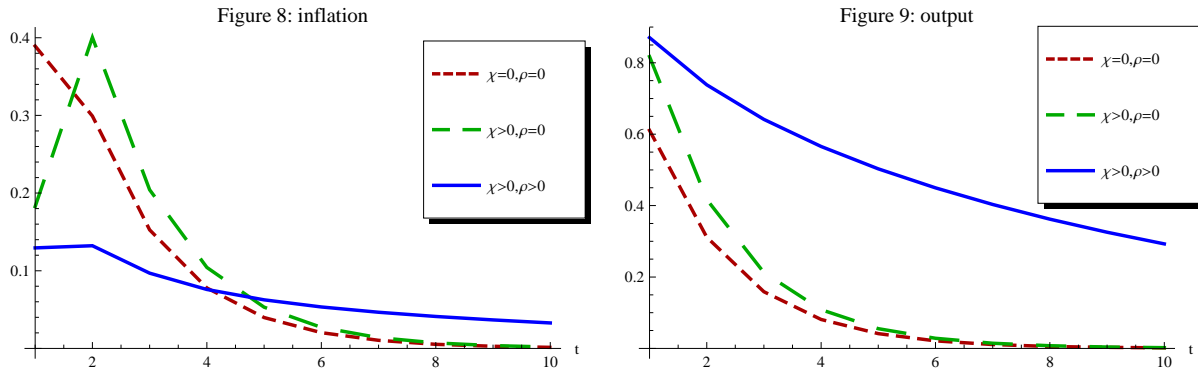
$$p_t = (1 - \lambda)p_{t-1} + \lambda P(p_{t-1}, \theta_t, \theta_{t-1}, \delta_{t-1}),$$

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 three properties. First, under this scenario the econometrician predicts that firms react to δ_t not because their signals are biased, but only because they believe that other firms will do so. Second, the two scenarios deliver similar qualitative properties as long as $\alpha > 0$. And third, the quantitative difference between the two scenarios vanishes as $\alpha \rightarrow 1$. 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 as in Proposition 4. We can then use this 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. 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. Moreover, this insight is not specific to the contemporaneous response of beliefs to the shock, but also to the entire dynamic adjustment of the beliefs. Indeed, in work that we do not report here because of space limitations, we have obtained similar results in a heterogeneous-priors variant of Woodford (2003), in which the shock never becomes common knowledge, thus allowing both first- and high-order beliefs converge to the truth only slowly over time. But whereas in Woodford (2003) the rate of convergence of higher-order beliefs is tightly connected to that of first-order beliefs, heterogeneous priors permits us to break this connection, so

that higher-order beliefs and prices may converge very slowly to their complete-information values even if first-order beliefs converge very fast.

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.¹¹

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, namely nominal shocks. In this section, building on Angeletos and La'O (2009b), we highlight how higher-order beliefs can *themselves* be the source of a certain type of fluctuations in the price level and real output—fluctuations that resemble the ones generated by cost-push, or mark-up, shocks.

For this purpose, we modify the model of the previous section as follows. 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.

¹¹The similarity we obtain between our two models is reminiscent of a point made in Lipman (2003). That paper shows that, if one fixes a specific hierarchy of beliefs, one cannot tell apart a common prior from heterogeneous priors from properties of any finite order of beliefs. This in turn suggests that in certain cases it may be possible to replicate, or approximate, the equilibrium behavior that obtains for any particular hierarchy of beliefs with either a common prior or heterogeneous priors. At the same time, because the reduced-form game that characterizes the general equilibrium of our economy admits a unique rationalizable outcome, and hence also a unique correlated equilibrium, the results of Kajii and Morris (1997) suggest that the complete-information equilibrium is robust to the introduction of incomplete information. It then follows that, with a common prior, one needs sufficient "noise" to make the predictions of the model under incomplete information sufficiently different.



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*

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

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 infrequent price adjustment may dampen the response of prices to nominal shocks. We did so through a series of variant models which progressively shifted focus from the stickiness of prices and the precision of available information about the shocks to the dynamics of higher-order expectations. We thus sought to highlight that quantifying the degree of price stickiness and the speed of learning (i.e., the rate at which first-order beliefs 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 also illustrated how the distinct role of higher-order beliefs could be readily accommodated within the Calvo framework without any sacrifice in analytical tractability—which in turn may pave the way to bringing the ideas of this paper closer to the data.

We thus hope that more effort will be devoted to quantifying the behavior of expectations and their implications for the degree of price inertia at the macro level. How can this be done? One possibility, at least in principle, would be to start conducting surveys of the higher-order beliefs of economic agents. We find this both impractical and unnecessary. As we further argue in Angeletos and La'O (2009a), in our eyes dispersed information—and higher-order beliefs—are merely modeling devices for capturing the uncertainty that economic agents may face about aggregate economic activity beyond the one that they face about the underlying fundamentals. Indeed, in macro models and games alike, higher-order beliefs matter only to the extent that they impact forecasts of the equilibrium actions of other agents. Furthermore, in most macro models this is typically summarized in forecasts of few macroeconomic variables, such as the price level and the level of aggregate output or employment. Therefore, we do not think that it is essential to collect data on the details of higher-order beliefs. Rather, we believe that data on forecasts of economic activity can provide more direct guidance for quantifying the type of effects we document here.

Turning to potential implications for monetary policy, we wish to make the following points. If one interprets the exogenous shock in our analysis as an innovation in monetary policy, one may conclude that our results justify strong real effects for exogenous changes to monetary policy. We would not necessarily favor such an interpretation. The theory we have presented here does not imply the same price inertia with respect to all shocks in the economy. Rather, it is crucial that there is non-trivial lack of common knowledge about the shock under consideration. But then note that information about changes in monetary policy is readily available and closely followed, not only from participants in financial markets but also from the general public when it seems to matter. Moreover, it is commonly understood that this is the case; there is a lot of communication in the market regarding monetary policy; and financial prices adjust within seconds to changes in monetary policy. In our eyes, these are indications that assuming common knowledge about the innovations to monetary policy might not be a terrible benchmark after all.

At the same time, we suspect a significant lack of common knowledge for a variety of other "structural" shocks hitting the economy, such as aggregate productivity shocks, financial shocks, labor-

market shocks, and so on. We then expect incomplete information to have more bite on monetary policy in a different dimension: the interest-rate rule followed by the central bank can affect how much the incompleteness of information about these shocks impacts equilibrium outcome. In particular, the central bank can control the degree of strategic complementarity in pricing decision by designing its response to realized inflation and output; in so doing, it can also mitigate the inertia effects we have documented here. Further exploring this possibility is left for future work.¹²

¹²The insight that the response of policy to macroeconomic outcomes can impact the decentralized use of information, and thereby the response of the economy to both the fundamentals, and noise draws from Angeletos and Pavan (2009). See also Angeletos and La'O (2008) and Lorenzoni (2009) for some recent work on the design of optimal monetary policy when information is dispersed.

Appendix

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

$$f(c_1, \lambda, \beta) = \alpha, \quad (16)$$

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,

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

It follows 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 established 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:

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

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

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

Clearly, c_1 and b_1 continue to be determined by the same equations as in the baseline model. Once again, we 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 Section 5 and that of Section 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 $(p_{t-1}, x_t, \theta_{t-1} \delta_{t-1}, \omega_t)$:

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

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

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

where

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

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

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

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:

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

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 (2008), “Dispersed Information over the Business Cycle: Optimal Fiscal and Monetary Policy,” MIT mimeo.
- [3] Angeletos, George-Marios, and Jennifer La’O (2009a), “Noisy Business Cycles,” paper prepared for the 2009 NBER Macroeconomics Annual.
- [4] Angeletos, George-Marios, and Jennifer La’O (2009b), “Sentiments,” MIT mimeo.
- [5] Angeletos, George-Marios, Guido Lorenzoni, and Alessandro Pavan (2008), “Wall Street and Silicon Valley: a Delicate Interaction,” MIT/Northwestern University mimeo.
- [6] Angeletos, George-Marios, and Alessandro Pavan (2007), “Efficient Use of Information and Social Value of Information,” *Econometrica* 75:4, 1103-1142.
- [7] Angeletos, George-Marios, and Alessandro Pavan (2009), “Policy with Dispersed Information,” *Journal of the European Economic Association* 7, 1-50
- [8] Bacchetta, Philippe, and Eric van Wincoop (2005), “Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?,” *American Economic Review* 95.
- [9] Barro, Robert (1976), “Rational Expectations and the Role of Monetary Policy,” *Journal of Monetary Economics* 2, 1-32.
- [10] Bils, Mark, and Pete Klenow (2004), “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy* 112, 947-985.
- [11] Calvo, Guillermo (1983), “Staggered Prices in a Utility Maximizing Framework,” *Journal of Monetary Economics* 12, 383-398.
- [12] Collard, Fabrice, and Harris Dellas (2005), “Misperceived vs Unanticipated Money: A Synthesis,” CNRS-GREMAQ/University of Bern mimeo.
- [13] Hellwig, Christian (2005), “Heterogeneous Information and the Welfare Effects of Public Information Disclosures,” UCLA mimeo.

- [14] Kajii, Atsushi, and Stephen Morris (1997), "The Robustness of Equilibria to Incomplete Information," *Econometrica* 65, 1283-1309.
- [15] Klenow, Pete, and Jonathan Willis (2007), "Sticky Information and Sticky Prices," *Journal of Monetary Economics* 54, 79-99.
- [16] Lipman, Barton (2003), "Finite Order Implications of Common Priors," *Econometrica* 71, 1255-1267.
- [17] Lorenzoni, Guido (2008), "A Theory of Demand Shocks," *American Economic Review*, forthcoming.
- [18] Lorenzoni, Guido (2009), "Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information," *Review of Economic Studies*, forthcoming.
- [19] Lucas, Robert E., Jr. (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4, 103-124.
- [20] Lucas, Robert E., Jr. (1975), "An Equilibrium Model of the Business Cycle," *Journal of Political Economy* 83, 1113-1144.
- [21] Mackowiak, Bartosz, and Mirko Wiederholt (2008), "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, forthcoming.
- [22] Mackowiak, Bartosz, and Mirko Wiederholt (2009), "Business Cycle Dynamics under Rational Inattention," European Central Bank/Nothwestern University mimeo.
- [23] 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.
- [24] Morris, Stephen, and Hyun Song Shin (2002), "The Social Value of Public Information," *American Economic Review* 92, 1521-1534.
- [25] Morris, Stephen, and Hyun Song Shin (2006), "The Inertia of Forward-Looking Expectations," *American Economic Review* 96, 1521-1534.
- [26] Nimark, Kristoffer (2007), "Dynamic Higher Order Expectations," Reserve Bank of Australia mimeo.

- [27] Nimark, Kristoffer (2008), "Dynamic Pricing and Imperfect Common Knowledge," *Journal of Monetary Economics* 55, 365-382.
- [28] 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.
- [29] Reis, Ricardo (2006), "Inattentive Producers," *Review of Economic Studies* 73, 793-821.
- [30] Sims, Christopher (2003), "Implications of Rational Inattention," *Journal of Monetary Economics* 50, 665-690.
- [31] Townsend, Robert M. (1983), "Forecasting the Forecasts of Others," *Journal of Political Economy* 91, 546-588.
- [32] 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.
- [33] Woodford, Michael (2008), "Information-Constrained State-Dependent Pricing," Columbia University mimeo.