

BIASED INFLATION FORECASTS

Hassan Afrouzi and Laura Veldkamp *
Columbia University

Abstract

Recent work finds that people’s beliefs about inflation are systematically upward biased. Since inflation expectations are central to the efficacy of monetary policy, understanding these expectations, and their biases, is important for policy. While one can always find preference-based explanations for bias, the fact that more informed agents have less upward bias, suggests some connection to information, as opposed to preferences. This paper proposes a rational Bayesian explanation for the bias: Agents with parameter uncertainty over positively-skewed distributions have a positive bias in their forecast. We use inflation and survey data to show that this mechanism can quantitatively explain the magnitude of the bias. The model implies that communicating about inflation skewness may be an important dimension of forward guidance.

1 Introduction

Inflation expectations are a ubiquitous ingredient in the economic decisions of households and firms. Whether a household is choosing how much to save, or a firm is deciding what price to set, a sound forecast about future inflation is essential for taking the optimal action. Understanding inflation expectations is particularly central to the efficacy of monetary policy. Monetary policy makers continuously refer to inflation expectations as a major driver of inflation and craft policies to target these expectations.

At the same time, inflation expectations are puzzling. Evidence from the U.S. and New Zealand shows that average inflation expectations of households and firms are systematically upward biased, by many percent. When asked to assign probabilities to different inflation outcomes, these agents exhibit large degrees of uncertainty and skewness, towards high inflation outcomes. Moreover, those who are more uncertain about inflation tend to have a more biased inflation forecast. The New Zealand case is particularly surprising because it has successfully implemented inflation targeting, for almost 3 decades. [Kumar, Afrouzi, Coibion](#)

*Please send comments to ha2475@columbia.edu, lv2405@columbia.edu.

and Gorodnichenko (2015) show that while both the official target of the Reserve Bank of New Zealand as well as the average inflation in New Zealand was 2% and the maximum in any quarter, since inflation targeting started, was 5%, about half of all firms forecast 1-year-ahead inflation to be 5% or higher (Figure V). Using evidence from Michigan Survey of Consumers and New York Fed’s Survey of Consumer Expectations, they also show the same patterns are observed for households in the U.S. (Figure III). Evidence about firms’ expectations in the U.S. is scarce, but in Bryan, Meyer and Parker (2015), U.S. firms exhibit similar upward bias when asked about the change in the overall prices in the economy: the average firm in their sample believed that yearly inflation was going to be 4.4% in September of 2015, when in reality inflation turned out to be less than 1%, far below the 2% target of the Federal Reserve. While forecasts always have error, an unbiased forecaster with decades of data should not make such large systematic errors.

Our contribution is to offer a belief-based explanation for this bias that can also explain its heterogeneity in the cross-section and over forecast horizons. We argue that forecast bias arises from rational Bayesian beliefs when the forecaster is uncertain about the parameters of the distribution from which inflation is drawn, but correctly believes that distribution to be positively skewed. Our two key assumptions – perceived skewness and subjective uncertainty about inflation probabilities – are both clearly reflected in the survey data. In the survey data from New Zealand, firms are asked to assign probabilities to different inflation outcomes which allows us to measure higher moments of their subjective beliefs. Figure (VI) plots the distribution of the standard deviation of these beliefs across firms. The average firm is highly uncertain about inflation, much more than the empirical distribution of inflation would seem to support. Figure (IV) shows the average probabilities that firms assign to each outcome. The distribution is skewed and firms assign large probabilities to high inflation outcomes.

There are some existing preference-based theories of forecast bias. These theories rely on asymmetries in preferences (Capistrán and Timmermann, 2009), robust control (Bhandari et al., 2017) or ambiguity aversion (Baqae, 2016). To distinguish between preference and belief-based explanations, we offer additional evidence about forecast bias. One additional piece of evidence is that professional forecasters and households/firms systematically have very different levels of bias. Of course, one can assume that their preferences differ in just the right way so as to explain this cross sectional difference. But our theory delivers these differences endogenously when professional forecasters observe more macro data and firms observe more marginal cost data. We show that the incentives of each type of agent to acquire information support this explanation. Second, we document that forecast bias at one-year and five-year horizons flips signs. This finding is particularly challenging to preference explanations. To explain it, the same agent, at the same moment in time, would have to have

a different objective or different concern for model misspecification at different horizons. For the belief-based explanation, this outcome is simply a result of different skewness at different frequencies, a feature observed in the data. Because policy makes inflation mean-reverting, the extreme positive inflation outcomes observed at one year horizons are not such outliers when averages with the surrounding 5 years of data. When the positive skewness in 5-year inflation data disappears, so does the agents' forecast bias.

But this simple Bayesian explanation raises further questions. First and foremost, if firms believe inflation is high, shouldn't they raise prices accordingly, which, in turn, would rationalize the high inflation belief? Our framework connects with the model in [Afrouzi \(2017\)](#) which explains this disconnect by showing that firms always prefer information about sector-specific prices, without paying separate attention to aggregates versus idiosyncratic shocks. We depart from this paper by focusing on how parameter uncertainty causes biased expectations, but use the findings of that paper to rationalize why prices of firms do not inherit the biases in their expectations: acquiring information about sector-specific prices reduces the firms parameter uncertainty and reduces their forecast bias for prices, in their sector. Because firms' forecasts of prices in their sector are unbiased, the prices they set to compete in their sector do not reflect their high aggregate inflation expectations. This is in contrast to [Maćkowiak and Wiederholt \(2009\)](#) where firms see separate signals about inflation and sector-specific shocks. In that setting since firms have to form their expectations about each shock separately, the bias in inflation expectations would pass-through to price one-to-one.

[Afrouzi \(2017\)](#) also documents that firms are far less uncertain about their own industry prices than they are about aggregate inflation; moreover, they are much less biased about their own industry prices than they are about aggregate inflation (Figure VIII). This is consistent with the prediction of our model that higher uncertainty should translate to larger biases.

Lastly, we draw several other testable predictions and verify them in the data: our model predicts that bias in inflation expectations should be countercyclical, whereas disagreement among agents should be procyclical, both of which are what we observe in the data for expectations of households in the Michigan Survey of Consumers.¹ Another puzzling observation in the Survey data is that households and firms revise their forecasts by a lot in short periods of time, an observation that is usually attributed to measurement error. After all, rational Bayesian agents who follow a stable process like inflation and know the parameters of the model should not revise their forecasts by much – which is the case for professional forecasters but not households or firms. Our model, however, makes sense of this observation

¹The robust control model of [\(Bhandari et al., 2017\)](#) also predicts a counter cyclical bias.

by relying on parameter uncertainty. Households and firms are far more uncertain about the parameters of the inflation process – because in our model they endogenously choose to see noisier data about inflation – which causes them to shift their estimates of tail event in light of new data and revise their forecasts by a lot.

Our message is that understanding forecasts requires relaxing the full-information assumptions of rational expectations econometrics. In such a full-information world, agents are assumed to know what the true distribution of economic outcomes is. Their only uncertainty is about what realization will be drawn from a known distribution. To model such a full-information forecaster, it makes sense to estimate a model on as much data as possible, take the parameters as given, and estimate the conditional expectations. But in reality, the macroeconomy is not governed by a simple, known model and we surely do not know its parameters. Instead, our forecast data suggests that forecasters estimate simple models to approximate complex processes and constantly use new data to update beliefs. Forecasters are not irrational. They simply do not know the economy’s true data-generating process. In such a setting, forecast means and sample means can behave quite differently. Our findings teach us that one important determinant of inflation expectations may be what is known about the distribution of economic outcomes.

Related Literature We are motivated by a large body of empirical evidence on inflation expectations of firms and households. It is well established that inflation expectations of households in the Michigan Survey of Consumers are systematically upward biased (Croushore, 1998; Carroll, 2003). Similar patterns have been documented in New York Fed’s Survey of Consumer Expectations (Kumar et al., 2015; Bhandari et al., 2017). The former is also our source for data on firms’ subjective beliefs. Coibion et al. (2018) use a more developed version of this data to go deeper in investigating how firms form their expectations and how these expectations affect their actions. Instead, we focus on the relationship between the bias and higher moments of firms’ beliefs, and use this data to provide evidence for our mechanism.

Our approach is inspired by two preceding papers that both estimate Bayesian forecasting models to describe agents’ beliefs. Cogley and Sargent (2005) use such a model to understand the behavior of monetary policy, while Johannes et al. (forthcoming) estimate a model of consumption growth to capture properties of asset prices. While the concept is similar, our use of a non-normal model, our measurement of uncertainty, and our study of the effects of learning about skewness distinguish our work from its predecessors. This approach is also motivated by Chen et al. (2013), which critiques models that give agents knowledge of parameters that econometricians cannot identify and Hansen’s Ely Lecture (Hansen, 2007)

which advocates modeling agents who struggle with real-time model selection.

2 Bayesian Belief Formation

To build up intuition for how parameter learning creates forecast bias, we start with the forecasting problem alone. Our forecasting model is one of the simplest forecasting models, an Kalman filter system, but with an exponential twist. The parameters of the exponential twist regulate the conditional skewness of outcomes. Because skewness is hard to learn in small samples, learning does not converge quickly. We prove the the combination of parameter uncertainty and skewness generates forecast bias. That is the key mechanism in the paper. The economic model that follows allows us to quantify the bias and answer economic questions about its consequences.

Before we describe the belief updating models, we start by defining terms. A model, denoted \mathcal{M} , has a vector a parameters θ . Together, \mathcal{M} and θ determine a probability distribution over a sequence of outcomes π_t . Let $\pi^t \equiv \{\pi_\tau\}_{\tau=1}^t$ denote a series of data (in our exercises, the inflation rates) available to the forecaster at time t . In every model, agent i 's information set $\mathcal{I}_{i,t}$ will include the model \mathcal{M} and the history π^t of observations up to and including time t . The state S_t , innovations, and the parameters θ are never observed.

Specifically, at each date t , the agent conditions on his information set $\mathcal{I}_{i,t}$ and forms beliefs about the distribution of π_{t+1} . We call the expected value $E(\pi_{t+1}|\mathcal{I}_{i,t})$ an agent i 's *forecast* and the square root of the conditional variance $Var(\pi_{t+1}|\mathcal{I}_{i,t})$ is what we call *subjective uncertainty*. Forecasters' forecasts will differ from the realized inflation rate. This difference is what we call a *forecast error*.

Definition 1. *An agent i 's forecast error is the distance, in absolute value, between the forecast and the realized inflation rate: $FE_{i,t+1} = |\pi_{t+1} - E[\pi_{t+1}|\mathcal{I}_{i,t}]|$.*

The average forecast error is then given by

$$\bar{FE}_{t+1} \equiv \int_0^1 FE_{i,t+1} di. \tag{1}$$

While we have only defined forecast errors and uncertainty for one-period-ahead forecasts, they can be similarly defined for other horizons.

Definition 2. *Uncertainty is the standard deviation of the time- $(t+1)$ inflation, conditional on an agent's time- t information: $U_{it} = \sqrt{E[(\pi_{t+1} - E[\pi_{t+1}|\mathcal{I}_{i,t}])^2|\mathcal{I}_{i,t}]}$.*

It is useful to observe that using definition (1) one could write uncertainty as the squared root of expected squared forecast error; $U_{it} = \sqrt{E[FE_{i,t+1}^2|\mathcal{I}_{i,t}]}$.

A Skewed Forecasting Model Figure (II) shows the distribution of inflation over time in New Zealand, and figure (IV) shows the average subjective belief distribution across firms. The evident skewness of these distributions suggest a departure from normality as positive skewness is a salient feature of inflation data.

Updating non-normal variables is typically cumbersome. Combining this with parameter uncertainty typically requires particle filtering, which is possible, but slow. We make this problem tractable by doing a change of measure. The Radon-Nikodym theorem tells us that, for any measure g that is absolutely continuous with respect to a measure induced by a normal distribution, we can find a change-of-measure function f such that $g(x) = \int f(x)d\Phi(x)$, where Φ is a normal CDF. If we estimate such an f function, then we use f^{-1} to take data from a skewed distribution and transform it into normal data, which simplifies the updating problem significantly.

Of course, allowing a forecaster to explore the whole function space of possible f 's is not viable. We focus the problem by considering a family of functions and allowing the forecaster to consider parameter estimates that govern the properties of the distribution. The change of measure function should have three desirable properties: 1) Its range is the real line; 2) it is monotone; and 3) it can be either globally concave or globally convex, depending on the estimated parameters. A class of transformations that satisfy this criteria is

$$f(\tilde{X}_t) = c + b \exp(-\tilde{X}_t) \tag{2}$$

If we have estimates for b and c , we can do a change of variable: Use $f^{-1}(\pi_t)$ to transform inflation into a variable $\tilde{X}_t = S_t + \sigma\epsilon_t$, which is a normally-distributed continuous variable with a persistent hidden state. Then we can write our skewed forecasting model as

$$\begin{aligned} \pi_t &= c + b \exp(-S_t - \sigma\epsilon_t) \\ S_t &= \rho S_{t-1} + \sigma_s \varepsilon_t \end{aligned}$$

where both ϵ_t and ε_t are standard normals. This change-of-variable procedure allows our forecaster to consider a family of non-normal distributions of inflation and convert each one into a linear-normal filtering problem with unknown parameters that can be estimated jointly using the same tools as in the previous section. The only additional complication is that the parameters b and c also need to be estimated.

We start with priors (see Appendix) and use MCMC techniques to form beliefs about the vector of parameters $\theta \equiv (\rho, \sigma, \sigma_s, c)$. For each parameter draw θ_i from the MCMC algorithm, we compute $E[\pi_{t+1}|\mathcal{I}_t, \theta_i]$ and $E[\pi_{t+1}^2|\mathcal{I}_t, \theta_i]$. We average these expectations over

all parameter draws and compute uncertainty as

$$U_t = \sqrt{E[\pi_{t+1}^2 | \mathcal{I}_t] - (E[\pi_{t+1} | \mathcal{I}_t])^2}.$$

The key feature of this model is that it produces a skewed distribution of outcomes and that the forecaster has to estimate parameters that govern the skewness. Skewness in this model is most sensitive to the b and c parameters because they govern the curvature of the transformation f of the normal variable. Any function with similar curvature, such as a polynomial or sine function, would deliver a similar mechanism.

2.1 Why Skewness and Model Uncertainty Generate Upward Bias in Forecasts?

Inflation forecasts are puzzling because it cannot be that over 70 years of post-war history, forecasters have not figured out that the sample mean is only half as high as their forecasts, on average. Our next result shows that these high forecasts are entirely rational for a Bayesian who believes that outcomes are positively skewed and faces parameter uncertainty.

Proposition 1. *Suppose that π is a random variable with a probability density function f that can be expressed as $f(y|\mu, \sigma) = \phi((g^{-1}(\pi) - \mu)/\sigma)$ where ϕ is a standard normal density and g is a convex function. Let the mean of π be $\bar{\pi} \equiv \int \pi f(\pi|\mu, \sigma) d\pi$. A forecaster does not know the true parameters μ and σ , but estimates probability densities $h(\mu)$ and $k(\sigma)$, with means μ and σ . The forecaster uses these parameter densities to construct a forecast: $\hat{\pi} \equiv \int \int \int \pi f(\pi|\mu', \sigma') h(\mu') k(\sigma') d\pi d\mu' d\sigma'$. Then $\hat{\pi} > \bar{\pi}$.*

The logic of the result is the following: If inflation is a convex transformation of a normal underlying variable, Jensen's inequality tells us that expected values will be systematically higher than the average realization. But by itself, Jensen's inequality does not explain the forecast bias because the expected inflation and the mean inflation should both be raised by the convex transformation (see figure I, left panel).

It must be that there is some additional uncertainty in expectations, making the Jensen inequality effect larger for forecasts than it is for the unconditional mean of the true distribution (see figure I, right panel). This would explain why our results tell us that most of the time the sample mean is less than the average forecast. If the agent knew the true parameters, he would have less uncertainty about π_{t+1} . Less uncertainty would make the Jensen effect smaller and lower his estimate of π_{t+1} , on average. Thus, it is the combination of parameter uncertainty and a skewed distribution that can explain the forecast bias.

3 Pricing Model

Next, we embed this Bayesian learning mechanism in a standard attention-allocation model of information choice and price setting. The purpose is to quantify the effect and to understand how expectations can be biased, while the prices firms set remain low.

There is a continuum of firms in the economy, indexed by $i \in [0, 1]$.

Profit function of firms. We assume that firm i 's profit at time t is given by the following profit function:

$$\Pi_{i,t} = \Pi(P_{i,t}, Q_t, Z_{i,t}) \quad (3)$$

where $P_{i,t}$ is the firm's own price, $Q_t \in \Theta_Q$ is the nominal aggregate demand (or money supply) and $Z_{i,t} \in \Theta_Z$ is an idiosyncratic shock to firm i 's profit. The function $\Pi(\cdot, \cdot, \cdot)$ is at least twice differentiable in all its arguments, and for any realization of $(Q_t, Z_{i,t})$, it is strictly concave and single peaked in its first argument. Moreover, the profit function is homogeneous of degree one in its first two arguments.²

Linear-Quadratic approximation. For $(Q, Z) \in \Theta_Q \times \Theta_Z$, let

$$P^*(Q, Z) \equiv \arg \max_{\hat{P}} \Pi(\hat{P}, Q, Z)$$

denote the price that maximizes the profit function. Let small letters denote log deviations of the corresponding variables from an arbitrary point $(\bar{P}, \bar{Q}, \bar{Z})$. Then, Taylor approximations around this point give:

$$p^*(q, z) = q + \alpha z + \mathcal{O}(\|q, z\|^2), \quad \alpha \equiv \frac{\bar{P} \partial_{13} \bar{\Pi}}{\bar{Z} \partial_{11} \bar{\Pi}} \quad (4)$$

$$\Pi(P, Q, Z) - \bar{\Pi} = -B(p - p^*(q, z))^2 + \mathcal{O}(\|q, z\|^3), \quad B \equiv -\bar{P}^2 \partial_{11} \bar{\Pi} > 0. \quad (5)$$

Equation (4) simply shows that up to a first order approximation the elasticity of firms' optimal price with respect to money supply is one – implied by the homogeneity of the profit function. Moreover, equation (5) shows that the firm's loss from charging a price $P \neq P^*(Q, Z)$ is proportional to the quadratic difference between the two.

²This is a standard profit function assumed in many macro models. For a microfoundation, see, for instance, (Woodford, 2003).

Prices with imperfect information. Let us assume that at time t , firm i 's information set is characterized by the set of all the signals that they have seen in the past:

$$\mathcal{I}_{i,t} \equiv \{s_{i,\tau} : \tau \leq t\}.$$

Given an information set at time t , the firm minimizes the quadratic loss in Equation (5):

$$L_{i,t} = B \min_{p_{i,t}} \mathbb{E} [(p_{i,t} - q_t - \alpha z_{i,t})^2 | \mathcal{I}_{i,t}].$$

The solution implies:

$$\begin{aligned} p_{i,t} &= \mathbb{E} [q_t + \alpha z_{i,t} | \mathcal{I}_{i,t}] \\ L_{i,t} &= \text{var}(q_t + \alpha z_{i,t} | \mathcal{I}_{i,t}) \end{aligned}$$

Optimality of one signal. Suppose in addition to q_t and $z_{i,t}$, the firms also learn about some parameters that we refer to by θ . Suppose for a given information set $\mathcal{I}_{i,t-1}$ at the beginning of t , the firm chooses a set of signals $s_{i,t}$ to solve:

$$\begin{aligned} \min_{s_{i,t}} \mathbb{E} [\text{var}(q_t + \alpha z_{i,t} | \mathcal{I}_{i,t}) | \mathcal{I}_{i,t-1}] + \phi \mathbb{I}(s_{i,t}; q_t, z_{i,t}, \theta | \mathcal{I}_{i,t-1}) \\ \text{s.t. } \mathcal{I}_{i,t} = \mathcal{I}_{i,t-1} \cup s_{i,t} \end{aligned}$$

where $\mathbb{I}(\cdot; \cdot | \cdot)$ is a conditional Shannon mutual information function – for instance Shannon – and $\phi > 0$ is the marginal cost of processing information.

Assumption 1. *The function $\mathbb{I}(\cdot; \cdot | \cdot)$ satisfies the following two properties:*

1. *if the two random variables X and Y are independent, then their mutual information is zero:*

$$X \perp Y \Rightarrow \mathbb{I}(X, Y) = 0.$$

2. *(Data Processing Inequality) For any Markov chain $X \rightarrow Y \rightarrow Z$ ($X \perp Z | Y$),*

$$\mathbb{I}(X, Y) \geq \mathbb{I}(X, Z).$$

Intuitively, the second condition means that if conditional on Y , Z reveals no information about X , then Y contains at least the same amount of information about X as Z .

Proposition 2. *Suppose the cost of information satisfies Assumption (1). Then, firms always prefer to see one signal at any point in time.*

Corollary 1. *In the limit where the marginal cost of information converges to zero, the optimal signal of firm i at time t converges to $s_{i,t} = q_t + \alpha z_{i,t}$ almost everywhere on the support of $q_t + \alpha z_{i,t}$ implied by $\mathcal{I}_{i,t-1}$.*

Equilibrium inflation and forecasts. Henceforth, we assume that the economy is in the limit where the marginal cost of processing information goes to zero, and all firms observe the optimal signal in this limit according to Corollary (1).

Define the aggregate price as the average price across firms, $p_t := \int_0^1 p_{i,t}$, and let $\pi_t := p_t - p_{t-1}$ denote the rate of inflation. Then in the limit:

$$p_t = \int_0^1 \mathbb{E}[q_t + \alpha z_{i,t} | \mathcal{I}_{i,t}] di = q_t. \quad (6)$$

The second part of the equality follows from Corollary (1) along with the fact that $z_{i,t}$'s are i.i.d. across firms. Hence, inflation is simply given by the growth in nominal aggregate demand (or money supply if velocity of money is constant):

$$\pi_t = \Delta q_t.$$

Finally, let $\mathbb{E}_{i,t}[\pi_{t+h}]$ denote the inflation forecast of firm i at time t for horizon h . Then, this forecast is given by

$$\mathbb{E}_{i,t}[\pi_{t+h}] = \mathbb{E}[\Delta q_{t+h} | \mathcal{I}_{i,t}], \quad \mathcal{I}_{i,t} = \{q_s + \alpha z_{i,s}\}_{s \leq t}. \quad (7)$$

3.1 How Can Expectations and Prices Be Disconnected?

There seems to be a disconnect between expectations and actions. In particular, if the average firm thinks that inflation was 4%, why did they only change their price by 2%? The results that follow rationalize this gap between expectations and actions.

Even though firms know their own optimal prices perfectly in the limit when the marginal cost of processing information goes to zero, they are still highly uncertain about the aggregate inflation:

$$\text{var}(\Delta q_t | \mathcal{I}_{i,t}) > \text{var}(\Delta q_t + \alpha \Delta z_{i,t} | \mathcal{I}_{i,t}) = 0.$$

Moreover, the inequality above implies that while q_t and $z_{i,t}$ are independent, firms'

optimal signal structure gives them a correlated posterior over these.

If the distribution of Δq_t (or alternatively $z_{i,t}$) is skewed, then firms' subjective distributions will inherit this property. Moreover, since firms choose their information structure such that they are not well-informed about inflation, their large uncertainty will translate to biased expectations about $\pi_t = \Delta q_t$.

3.2 How Should Inflation Bias Differ Between Forecasters and Firms?

The following result show that forecasters have small bias because the abundant information they observe about inflation causes them to have little parameter uncertainty. In contrast, firms observe marginal cost information, which is a noisier source of inflation data. Because their data is noisier, it is less informative about the parameters of the inflation distribution. The substantial parameter uncertainty, combined with positive inflation skewness, is what explains firms' positive bias in inflation expectations.

3.3 Inflation Bias at Longer Horizons

The data show substantially less upward inflation bias at long horizons. Here, we derive results on how Bayesian forecast bias should vary by horizon. We show that if inflation is sufficiently mean-reverting, then longer horizon bias should be lower because the mean-reversion makes multi-year average inflation statistics less skewed.

4 Quantitative Results: How Much Bias Can Inflation Skewness Explain?

Next, we use historical U.S. inflation data to estimate the forecasting model and determine how much bias it produces, in the aggregate, in the cross section, and across forecasting horizons.

Model estimation procedure To compute forecasts, we use Bayesian updating. A forecast is a conditional expectation of next-period inflation, where the expectation is taken over unknown parameters, states, and inflation realizations. Using the law of iterated expectations, we can write this forecast as:

$$E(\pi_{t+1}|\pi^t) = \int \int \int \pi_{t+1} p(\pi_{t+1}|\theta, S_{t+1}, \pi^t) p(S_{t+1}|\theta, \pi^t) p(\theta|\pi^t) d\theta dS_{t+1} d\pi_{t+1} \quad (8)$$

The first probability density function, $p(\pi_{t+1}|\theta, S_{t+1}, \pi^t)$, is the probability of $t+1$ inflation, given the state and the parameters.

The second probability density function, $p(S_{t+1}|\theta, \pi^t)$, is the probability of a hidden state in a Kalman filtering system. This is a (conditional) normal density. To see this, assume for a moment that the parameters are known. Then, if we have estimates for b and c , we can do a change of variable: Use $f^{-1}(\pi_t)$ to transform inflation into a variable $X_t = S_t + \sigma\epsilon_t$, which is a normally-distributed continuous variable with a persistent hidden state. This change-of-variable procedure allows our forecaster to consider a family of non-normal distributions of inflation and convert each one into a linear-normal (Kalman) filtering problem with unknown parameters that can be estimated jointly using the standard Bayesian estimation techniques. The following equations describe the conditional mean and variance of the first two probability terms, jointly

$$E[\pi_{t+1}|\pi^t, \theta, \mathcal{M}] = c + b \exp\left(-E[S_{t+1}|\pi^t, \theta, \mathcal{M}] + \frac{1}{2}\text{Var}[S_{t+1}|\pi^t, \theta, \mathcal{M}] + \frac{1}{2}\sigma^2\right)$$

where the following recursion characterizes the updating of state belief

$$E[S_t|\pi^t, \theta, \mathcal{M}] = (1 - K_t) E[S_t|\pi^{t-1}, \theta, \mathcal{M}] + K_t \ln((\pi_t - c)/b)$$

and where the term $K_t = \text{Var}[\ln((\pi_t - c)/b)|\pi^{t-1}, \theta, \mathcal{M}] (\text{Var}[\ln((\pi_t - c)/b)|\pi^{t-1}, \theta, \mathcal{M}] + \sigma_s^2)^{-1}$ is the Kalman gain. The conditional variance is given by

$$\text{Var}[\ln((\pi_t - c)/b)|\pi^t, \theta, \mathcal{M}] = \rho^2 \left[\frac{1}{\text{Var}[\ln((\pi_t - c)/b)|\pi^{t-1}, \theta, \mathcal{M}] + \frac{1}{\sigma_s^2}} \right]^{-1} + \sigma^2. \quad (9)$$

Finally, the third probability density function is the probability of the parameter vector θ , conditional on the t -history of observed data. To estimate the posterior parameters distribution, we employ Markov Chain Monte Carlo (MCMC) techniques. At each date t , the MCMC algorithm produces a sample of parameter vectors, $\{\theta^d\}_{d=1}^D$, such that the probability of any parameter vector θ^d being in the sample is equal to the posterior probability of those parameters, $p(\theta^d|\pi^t)$. Therefore, we can compute an approximation to any integral by averaging over sample draws: $\int f(\theta)p(\theta|\pi^t)d\theta \approx 1/D \sum_d f(\theta^d)$.

To estimate uncertainty, we compute these probability density terms and integrate numerically to get a forecast. In similar fashion, we also calculate $E(\pi_{t+1}^2|\pi^t)$. Applying the variance formula $\text{Var}(\pi_{t+1}|\pi^t) = E(\pi_{t+1}^2|\pi^t) - E(\pi_{t+1}|\pi^t)^2$, and taking the square root yields uncertainty: $U_t = \sqrt{\text{Var}(\pi_{t+1}|\pi^t)}$.

Our forecaster needs prior distributions over all the parameters to start the updating

process. We start with a flat prior, estimate each parameter³ on inflation data from 1947:Q2-1968:Q3, and use the mean and variance of this estimate as the mean and variance of prior beliefs. (See appendix for more details and prior estimation results.) Starting in quarter 4 of 1968, each period, the agent observes π_t and revisions of previous quarters' data and updates his beliefs about future inflation using (8). We start the estimation of the model in 1968:Q4 because this is the first quarter for which we have forecasts from the Survey of Professional Forecasters.

Properties of Model Forecast Bias

5 Empirical Support:

We end by testing predictions of our hypothesis. The first prediction is that agents' perceived uncertainty should predict the amount of bias in their forecast. Since bias comes from parameter uncertainty and parameter uncertainty is one component of one's inflation uncertainty, these two should be positively related.

Figure (VII) shows the scatter plot of firms' subjective uncertainty vs. their forecasts. Firms that report higher uncertainty (blue dots higher on y-axis) have higher inflation forecasts (further to the right), on average. Table (I) shows that the relationship is statistically significant. This effect is also economically large. An 8% inflation expectation is associated with double the uncertainty about the inflation rate.

The second prediction is that firms are more certain about their own industry price. In our model, this result arises because firms learn about firm-specific, or industry-specific, information because that is a more efficient way to acquire the information needed to set an optimal price.

Figure (VIII) shows the size of firms' nowcast errors ($E[\pi_t|\mathcal{I}_t] - \pi_t$) about both inflation and their industry prices. The two sets of forecast errors are strikingly different. Firms have highly accurate, nearly unbiased beliefs about their own industry. They have highly biased and quite inaccurate beliefs about aggregate inflation.

This prediction is a significant one because it helps to distinguish information, which likely differs between professional forecasters and firms, from preferences, which might not

³In the results we present, we introduced one modification. Notice that the b parameter governs the mean of the X_t process. To see this, note that for $b < 0$, we can rewrite $b \exp(-X_t) = -\exp(-X_t + \ln(|b|))$. To streamline our code, we simply remove the time- t sample mean of the X_t and set $b = -1$. After estimating the parameters of the mean-zero process, we add back in the sample mean. This approach is supported by the fact that when we have estimated b in less complex settings, we come up with consistently negative values and quantitatively similar estimates.

vary across the population. Of course, one might say that agents with ambiguity about the correct model could learn from more data. Then professional forecasters might consider a smaller set of models and have lower errors. This explanation is hard to rule out, as is any other explanation about how ambiguity might vary in the cross-section or across environments or forecasts. In contrast, our effect is tightly disciplined by data. The amount of skewness in the data governs the size of the bias.

Furthermore, if the ambiguity effect is disciplined by Bayesian uncertainty, we are back to the question of why Bayesian uncertainty is associated with forecast bias. Ambiguity or preference twists are possible, but not necessary to explain this. One just needs to incorporate the empirically relevant amount of skewness in the data. Skewness is not an assumption, it is a measurable feature of data.

Another prediction of the theory is that belief bias should differ by forecast horizons, if skewness differs by forecast horizon. Next, we plan to test that prediction in the quantitative model and in the data.

6 Conclusions

Most approaches to understanding inflation expectations ignore parameter estimation uncertainty. Sometimes referred to as “rational expectations econometrics,” the traditional approach entails estimating a model on the full sample of data and then treating the estimated parameters as truth to infer what the optimal forecast was in each period in the past. In doing this, the econometrician is assuming that the uncertain agent knows the true distribution of outcomes at every moment in time and is only uncertain about which outcome will be chosen from this distribution. Assuming such precise knowledge of the economic model rules out most uncertainty and ignores many sources of uncertainty shocks.

We explore the properties of forecasts when an agent is not endowed with knowledge of the true economic model and needs to estimate it, just like an econometrician. When the agent considers skewed distributions of outcomes, new data or real-time revisions to existing data can change his beliefs about the skewness of the distribution, and thus the probability of extreme events. Small changes in the estimated skewness can increase or decrease the probability of these tail events many-fold. Because tail events are so far from the mean outcome, changes in their probability have a large effect on forecasts. Thus, our message is that beliefs about extreme events that are never observed, but whose probability is inferred from a forecasting model, are responsible for much of forecast bias. For inflation, the relevant extreme events are positive, leading to positive inflation bias.

Our mechanism could also help explain important economic phenomena such as debt

crises and inflation risk premia. In a model where firms have debt, changes in skewness estimates would translate into changes in risk premia. The model would then tell us what kinds of events trigger high default risk and debt crises.

References

- Afrouzi, Hassan**, “Strategic Inattention, Inflation Dynamics and the Non-Neutrality of Money,” 2017. Mimeo.
- Baqae, David**, “Asymmetric inflation expectations, downward rigidity of wages, and asymmetric business cycles,” 2016. Mimeo.
- Bhandari, Anmol, Jaroslav Borovicka, and Paul Ho**, “Survey data and subjective beliefs in business cycle models,” 2017. University of Minnesota Working Paper.
- Bryan, Michael F, Brent Meyer, and Nicholas Parker**, “The Inflation Expectations of Firms: What Do They Look Like, Are They Accurate, and Do They Matter?,” 2015. Working Paper No. 2014-27a.
- Capistrán, Carlos and Allan Timmermann**, “Disagreement and biases in inflation expectations,” *Journal of Money, Credit and Banking*, 2009, 41 (2-3), 365–396.
- Carroll, Christopher D**, “Macroeconomic expectations of households and professional forecasters,” *the Quarterly Journal of economics*, 2003, 118 (1), 269–298.
- Chen, Hui, Winston Dou, and Leonid Kogan**, “Measuring the Dark Matter in Asset Pricing Models,” 2013. MIT working paper.
- Cogley, Timothy and Thomas Sargent**, “The Conquest of US Inflation: Learning and Robustness to Model Uncertainty,” *Review of Economic Dynamics*, 2005, 8, 528–563.
- Coibion, Olivier and Yuriy Gorodnichenko**, “Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation,” *American Economic Journal: Macroeconomics*, 2015, 7 (1), 197–232.
- , – , and **Saten Kumar**, “How Do Firms Form Their Expectations? New Survey Evidence,” *American Economic Review*, 2018, 108 (9), 2671–2713.
- Croushore, Dean**, “Evaluating inflation forecasts,” 1998. Working Paper Number 98-14.
- Hansen, Lars**, “Beliefs, Doubts and Learning: Valuing Macroeconomic Risk,” *American Economic Review*, 2007, 97(2), 1–30.
- Johannes, Michael, Lars Lochstoer, and Yiqun Mou**, “Learning About Consumption Dynamics,” *Journal of Finance*, forthcoming.
- Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko**, “Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand,” *Brookings Papers on Economic Activity*, 2015, Fall, 151–208.
- Maćkowiak, Bartosz and Mirko Wiederholt**, “Optimal sticky prices under rational inattention,” *American Economic Review*, 2009, 99 (3), 769–803.
- Woodford, Michael**, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2003.

7 Figures

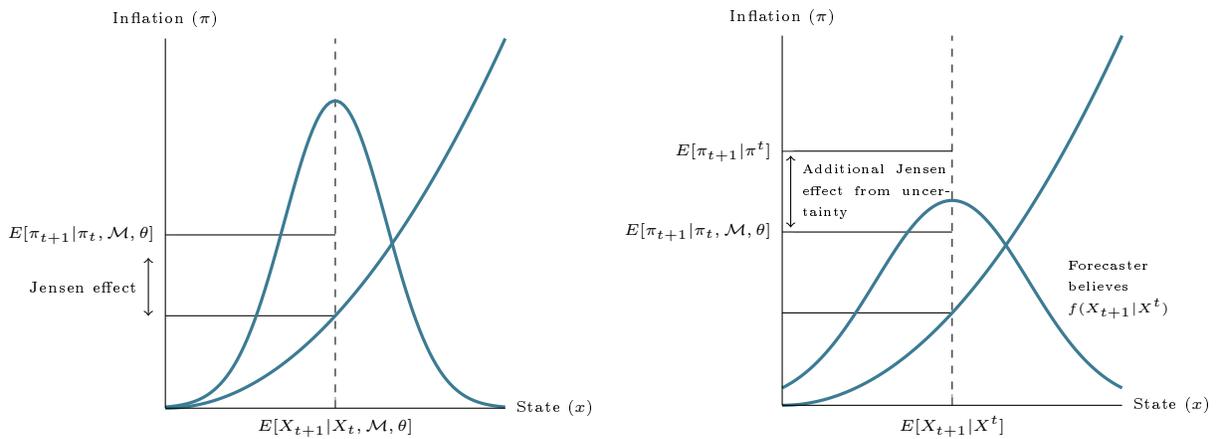


Figure I: Explaining why average forecasts are higher than mean inflation. The result has two key ingredients: The forecaster faces more uncertainty than he would if he knew the true distribution of outcomes, and a Jensen inequality effect from the convex change of measure.

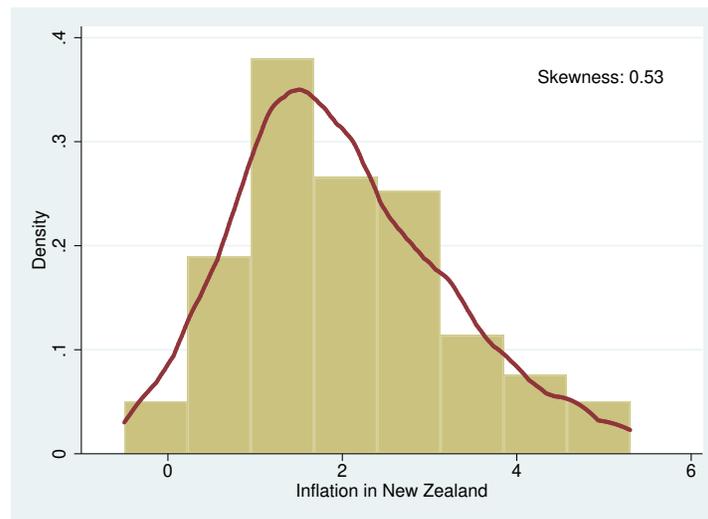


Figure II: The figure shows the distribution of realized annual inflation in New Zealand since December of 1991. The distribution is positively skewed. Source: Quarterly CPI inflation data released by the Reserve Bank of New Zealand.

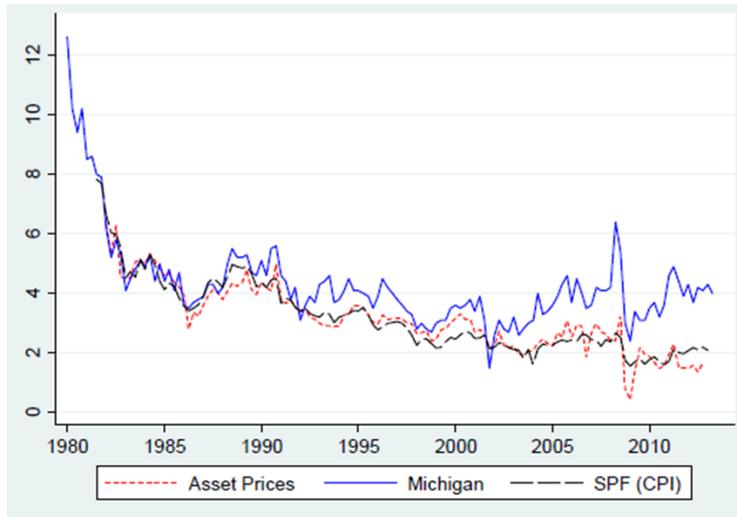


Figure III: The figure shows three different measures of inflation expectations. the blue line is the average one year ahead inflation expectation of households from the Michigan survey of consumers. The black line shows the one year ahead inflation expectation of professional forecasters, and the red line is the imputed measure of inflation expectations from asset prices. Households are consistently biased upwards in their expectations. Source: [Coibion and Gorodnichenko \(2015\)](#).

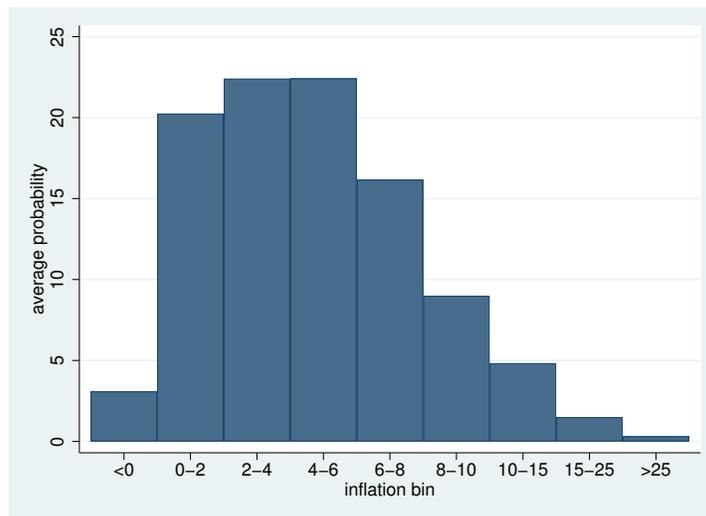


Figure IV: Average subjective distribution across firms. The figure shows the average probability assigned by firms to different inflation outcomes. Inflation has been anchored around 2% in New Zealand for the last 27 years. Nevertheless, firms assign high probabilities to large and unlikely inflation outcomes at the right tail. Source: [Kumar, Afrouzi, Coibion and Gorodnichenko \(2015\)](#).

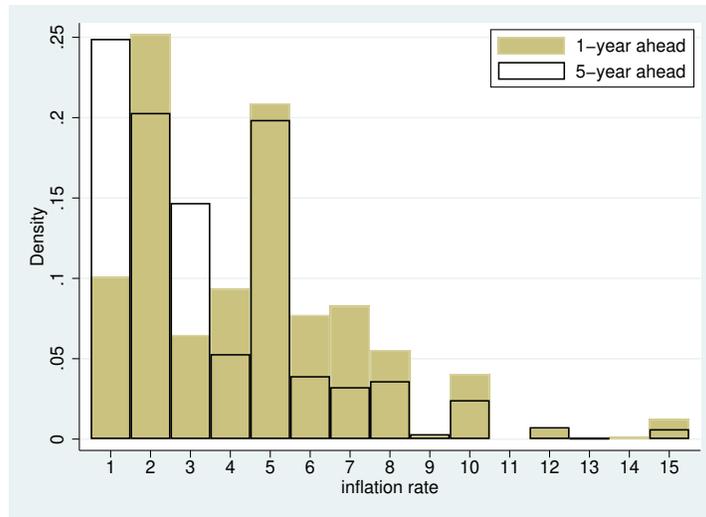


Figure V: The figure shows the distribution of 1 year ahead inflation forecasts across firms. In spite of low and stable inflation in New Zealand, firms disagree a lot about inflation and the average forecast is substantially higher than that of the Reserve Bank of New Zealand. Source: [Kumar, Afrouzi, Coibion and Gorodnichenko \(2015\)](#).

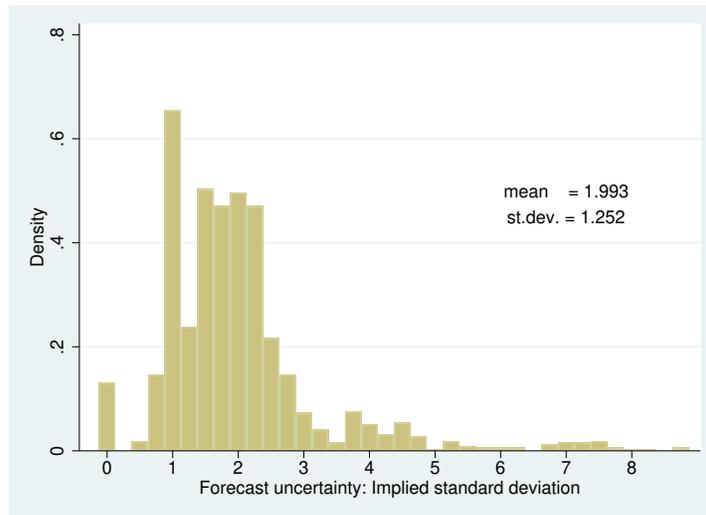


Figure VI: The figure shows the cross sectional distribution of subjective uncertainty of firms about inflation. Despite stable inflation in New Zealand, firms are highly uncertain about inflation. Source: [Kumar, Afrouzi, Coibion and Gorodnichenko \(2015\)](#).

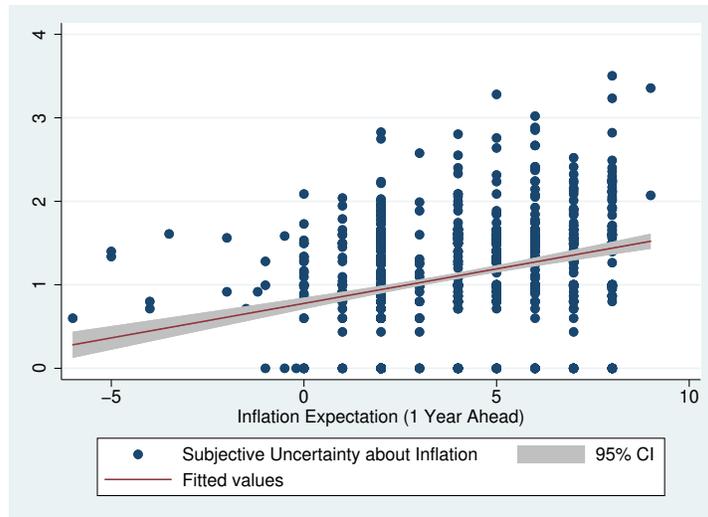


Figure VII: The figure shows the relationship between firms' inflation expectations and their subjective uncertainty about inflation, measured as the standard deviation of firms' reported distributions. More uncertain firms forecast inflation to be higher. Source: authors' analysis.

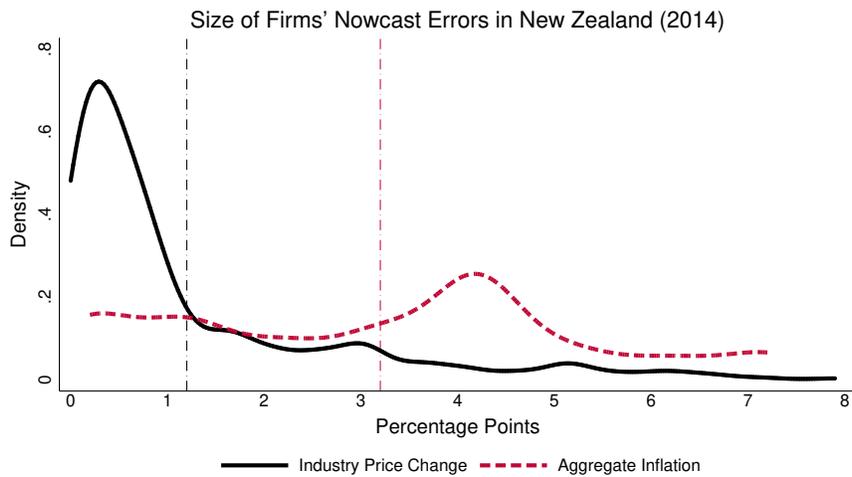


Figure VIII: The figure shows the average nowcast error of firms about their own industry price change versus aggregate inflation in the prior 12 months leading to the survey. Firms are less biased in recalling their own industry prices than aggregate inflation. Source: Afrouzi (2017).

8 Tables

	Subjective Uncertainty	
	(1)	(2)
Constant	0.752*** (0.044)	0.726*** (0.056)
Inflation Expectation (1 Year Ahead)	0.094*** (0.011)	0.092*** (0.012)
Industry FE	No	Yes
Observations	1032	903

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table I: The table shows the result of regressing firms' subjective uncertainty – measured as the standard deviation of their reported distribution about inflation – on their inflation forecasts. Firms with higher uncertainty forecast inflation to be higher.

Appendix A Proofs

Proof of Proposition (2). We prove this result by showing that for any choice $s_{i,t}$, there is a singleton choice that induces smaller losses to the firm. To see this, for any arbitrary $s_{i,t}$, let

$$\hat{s}_{i,t} \equiv \mathbb{E}[q_t + \alpha z_{i,t} | \mathcal{I}_{i,t-1} \cup s_{i,t}].$$

Notice that $\hat{s}_{i,t}$ is a singleton independent of how many signals there are in $s_{i,t}$. Moreover, notice that

$$\begin{aligned} & \mathbb{E}[\text{var}(q_t + \alpha z_{i,t} | \mathcal{I}_{i,t-1} \cup s_{i,t}) | \mathcal{I}_{i,t-1}] \\ &= \mathbb{E}[\mathbb{E}[(q_t + \alpha z_{i,t} - \hat{s}_{i,t})^2 | \mathcal{I}_{i,t-1} \cup s_{i,t}] | \mathcal{I}_{i,t-1}] \\ &= \mathbb{E}[\mathbb{E}[(q_t + \alpha z_{i,t} - \hat{s}_{i,t})^2 | \mathcal{I}_{i,t-1} \cup \hat{s}_{i,t}] | \mathcal{I}_{i,t-1}] \\ &= \mathbb{E}[\text{var}(q_t + \alpha z_{i,t} | \mathcal{I}_{i,t-1} \cup \hat{s}_{i,t}) | \mathcal{I}_{i,t-1}] \end{aligned} \quad (10)$$

Now, let $x_{i,t} := (q_t, z_{i,t}, \theta)$ represent the vector of all things that firm i learns about at time t . Note that $x_{i,t} | \mathcal{I}_{i,t-1} \rightarrow s_{i,t} | \mathcal{I}_{i,t-1} \rightarrow \hat{s}_{i,t} | \mathcal{I}_{i,t-1}$. Then, by data processing inequality:

$$\mathcal{I}(\hat{s}_{i,t}; x_{i,t} | \mathcal{I}_{i,t-1}) \leq \mathcal{I}(s_{i,t}; x_{i,t} | \mathcal{I}_{i,t-1}) \quad (11)$$

Combine equations (10) and (11) to observe that the firm prefers $\hat{s}_{i,t}$ to $s_{i,t}$. \square

Proof of Corollary (1). At any time t , for a given initial information set $\mathcal{I}_{i,t-1}$, and a sequence $\phi_n \rightarrow 0$, let $\{s_{i,t}^n\}_{n \in \mathbb{N}}$ be the sequence of solutions to:

$$\begin{aligned} s_{i,t}^n &\equiv \arg \min_{s_{i,t}} \mathbb{E} [\text{var}(q_t + \alpha z_{i,t} | S_i^t) | \mathcal{I}_{i,t-1}] + \phi_n \mathcal{I}(s_{i,t}; q_t, z_{i,t}, \delta_{i,t} | \mathcal{I}_{i,t-1}) \\ \text{s.t. } & \mathcal{I}_{i,t} = \mathcal{I}_{i,t-1} \cup s_{i,t} \end{aligned}$$

Note that any solution should have the property that

$$\lim_{n \rightarrow \infty} \text{var}(q_t + \alpha z_{i,t} | \mathcal{I}_{i,t-1} \cup s_{i,t}^n) = 0. \quad (12)$$

Otherwise, the loss of the firm is bounded below by some $\bar{l} > 0$ even though that the cost of information is going to zero. So we can construct another sequence of signals that are strictly better than the solution – a contradiction. Therefore, Now, given this sequence of solutions, let $\tilde{s}_{i,t}^n \equiv \mathbb{E}[q_t + \alpha z_{i,t} | \mathcal{I}_{i,t-1} \cup s_{i,t}^n], \forall n \in \mathbb{N}$. By the previous proposition we know that for all $n \in \mathbb{N}$ the firm prefers $\tilde{s}_{i,t}^n$ to $s_{i,t}^n$. So, $\{\tilde{s}_{i,t}^n\}_{n \in \mathbb{N}}$ is also a solution to the sequence of problems above. Finally, Equation (12) implies that almost everywhere

$$\lim_{n \rightarrow \infty} \mathbb{E}[q_t + \alpha z_{i,t} | \mathcal{I}_{i,t-1} \cup s_{i,t}^n] = q_t + \alpha z_{i,t} \quad (13)$$

Hence, $\tilde{s}_{i,t}^n \rightarrow_{a.e.} q_t + \alpha z_{i,t}$ \square