Uncertainty and Disagreement: Theory and Macroeconomic Applications*

Anand Bharadwaj †

January 4, 2025

Abstract

Disagreement across forecasters is commonly used as an empirical proxy for forecast uncertainty in macroeconomics, finance and other disciplines, but comparatively little is known about the theoretical soundness of this methodology. In this paper, we derive novel characterizations of and bounds for disagreement and uncertainty in common-prior Bayesian learning models with arbitrary information structures. We apply these results to study the dynamics of disagreement and uncertainty in canonical macroeconomic models of noisy and sticky information, with a focus on the comparative statics of these quantities with respect to the degree of informational frictions and other key parameters. We then relate the theoretical findings to empirical data on inflation and unemployment forecasts from the Survey of Professional Forecasters and illustrate the capacity of the disagreement/uncertainty ratio to discriminate between competing models of informational frictions.

Keywords: Uncertainty, disagreement, information frictions, expectations, survey forecasts
JEL Classification: D82, D84, E17, E37

*I am indebted to Yuriy Gorodnichenko, Martin Lettau, Emi Nakamura and Chris Shannon for their guidance on this paper. I am also especially grateful to Michael Woodford for many suggestions that have greatly improved the paper. For valuable comments and feedback, I also thank Robert Barro, Jamie Cross, Jonathon Hazell, Barry Hughes, Chen Lian, Demian Pouzo, Andreas Schaab, Benjamin Schoefer, Kannan Sethuraman, Michael Smith and Jón Steinsson.

†Department of Economics, University of California, Berkeley. Email: anand_bharadwaj@berkeley.edu.

1 Introduction

"Expectations matter," as Coibion and Gorodnichenko (2015) pithily put it. Much of modern macroeconomic research has been devoted to the development and empirical testing of models of expectations, ever since the seminal work of Lucas Jr (1972). More recently, three significant theoretical contributions in this space have been the noisy information models of Woodford (2003) and Sims (2003) and the sticky information model of Mankiw and Reis (2002). These authors, among many others, have explored the rich effects of informational frictions on macroeconomic dynamics and their implications for the optimal conduct of stabilization policy.

To test these theoretical hypotheses, a burgeoning literature in macroeconomics has leveraged survey data – both from households and firms, as well as professional forecasters – to shed light on how economic agents process information and form expectations. This literature, which includes Coibion and Gorodnichenko (2012), Coibion and Gorodnichenko (2015) and Bordalo et al. (2020), among a plethora of other contributions, has provided substantial novel insights into the nature and structure of expectations, how forecasters integrate new information into their forecasts, and, as a corollary, which macroeconomic models may be the most plausible.

In principle, forecasts should provide information about both expectations and perceived uncertainty. Many, but not all, macroeconomic surveys only request respondents to provide point predictions. In such scenarios, it is often important to develop indirect measures of forecast uncertainty, and a measure that has been commonly used in empirical practice is the cross-sectional dispersion, or disagreement, of point forecasts across survey respondents.

However, there is surprisingly little theoretical foundation for the use of this proxy. As Zarnowitz and Lambros (1987) suggest, the relationship between uncertainty and disagreement can be conceptualized within a two-by-two matrix. Intuitively, low uncertainty and low disagreement jointly occur when forecasters have precise signals and similar information sets. High uncertainty and low disagreement occur together when forecasters have similar

information sets but their signals and relatively uninformative. High uncertainty and high disagreement can occur either when forecasters possess diffuse information sets or when their priors are sufficiently heterogeneous. In some sense, the most atypical case is that of high disagreement but low uncertainty, which would typically be inconsistent with the common prior doctrine, under which differences in forecasts solely occur due to differences in information. A contribution of this paper is that we are able to formulate these ideas developed by Zarnowitz and Lambros (1987) in a theoretically rigorous manner that encompasses a wide class of standard models of informational frictions.

One important macroeconomic survey that collects information about both point and distribution forecasts is the Survey of Professional Forecasters (SPF). Several authors have leveraged data from the SPF to document stylized facts about the dynamics of uncertainty and disagreement, including Giordani and Söderlind (2003), Lahiri and Liu (2006), D'Amico and Orphanides (2008) and Lahiri and Sheng (2010). The literature has broadly found that disagreement and uncertainty are highly correlated, but there is still much contention about how to interpret this comovement and whether it is a theoretically robust phenomenon.

The use of disagreement as an empirical proxy for uncertainty is also common in other disciplines, including accounting (e.g. Barron and Stuerke, 1998; Barron et al., 1998); finance (e.g. Diether et al., 2002; Zhang, 2006) and operations management (e.g. Gaur et al., 2007). This literature has provided support for disagreement as a proxy for uncertainty by establishing empirical relationships between forecast errors and dispersion, but has devoted comparatively less attention to the theoretical foundations of this relationship.

In discussing the previous economics literature on disagreement and uncertainty, a few themes are evident. First, there is some degree of contention pertaining to the necessity or lack thereof for heterogeneity in priors to meaningfully capture the structure of forecasters' expectations. For example, Patton and Timmermann (2010) suggest that the observed characteristics of disagreement in the SPF cannot be rationalized without heterogeneity in priors or models and Farmer et al. (2024) find that heterogeneity in priors is a parsimonious

way to rationalize a broad range of forecast anomalies. Notwithstanding this, other authors, such as Coibion and Gorodnichenko (2015) have disputed the conclusion of Patton and Timmermann (2010), arguing that "heterogeneity about long-run means is not playing a quantitatively significant role in accounting for the predictability of inflation forecast errors."

In this paper, we do not take a strong positive stand on the "correct" way of modeling informational rigidities, but we invoke the common prior assumption throughout. In this sense, our results can be thought of as a valuable theoretical benchmark in the absence of the additional possibility that forecasters "agree to disagree". While some valuable theoretical contributions about the behavior of disagreement and uncertainty have been made, for example, by Patton and Timmermann (2010) and Lahiri and Sheng (2010), these papers have typically made restrictive assumptions about the statistical properties of forecasts, as opposed to the informational regimes that forecasters face. By allowing for a very general class of informational primitives, both in the presence and absence of dynamic state persistence, we obtain a richer array of theoretical results and challenge some of the conclusions highlighted in the previous literature.

We make four main contributions in this paper. First, we derive a general theorem that shows that expected disagreement is always less than or equal to expected uncertainty under a common prior with rational expectations. Second, we provide an analytical characterization of disagreement and uncertainty under different models of expectations (namely, noisy and sticky information) in both static and dynamic environments. Third, we derive a rich array of comparative statics for the ratio of disagreement to uncertainty with respect to the model primitives. Fourth, we relate these theoretical findings to empirical data from the Survey of Professional Forecasters to disambiguate between competing models of informational frictions.

An important novel aspect of this paper is the focus on the disagreement/uncertainty ratio. We show that this ratio provides valuable information about the nature of informational frictions that forecasters face, an idea that is thematically complementary to the

regression of forecast errors on forecast revisions developed in Coibion and Gorodnichenko (2015). Indeed, under both noisy and sticky information with a common prior, the ratio of expected disagreement to expected uncertainty cannot exceed unity (as predicted by the main theorem). However, under noisy information, the disagreement/uncertainty ratio is increasing in fundamental volatility (that is, the variance of the underlying state variable) but does not change in response to shocks. In contrast, under sticky information, the disagreement/uncertainty ratio is time and state-dependent, but taking the expectation with respect to shocks, its unconditional level is invariant with respect to fundamental volatility.

Taken together, since the empirical evidence shows that disagreement and uncertainty are strongly correlated, disagreement is relatively unresponsive to shocks (Coibion and Gorodnichenko, 2015), and the disagreement/uncertainty ratio seldom exceeds unity except during periods that appear to have high fundamental volatility (such as the Great Recession or the COVID-19 pandemic), we concur with Coibion and Gorodnichenko (2015) that the noisy information model fits the observed behavior of forecasters better than the sticky information model. As we discuss, letting fundamental volatility be stochastic in the noisy information model allows us to plausibly match the empirical time series of the disagreement/uncertainty ratio. However, there are some challenges in simultaneously explaining the seasonal behavior of fixed-event forecasts, as we discuss later in the paper. While these results are by no means a repudiation of other hypotheses that invoke heterogeneous priors or models, they do suggest the relative adequacy of the common prior assumption (Coibion and Gorodnichenko, 2015) and motivate further investigation of the disagreement/uncertainty ratio in future work.

2 Empirical Motivation

To motivate the theoretical analysis, we use data from the Survey of Professional Forecasters (SPF), which is a quarterly survey conducted by the Federal Reserve Bank of Philadelphia, where a panel of professional forecasters from Wall Street financial firms, banks, economic

consulting firms and large corporations are tasked with providing forecasts of a range of macroeconomic variables at different horizons. A unique feature of the SPF is that participants are asked not only to provide point forecasts, but also subjective probability estimates for some of these variables, that is, their perceived probabilities of these variables landing in specific ranges at a given horizon. The SPF is one of the few major macroeconomic surveys that provide subjective probability distributions, making it an ideal choice for this analysis.

However, some data limitations must be acknowledged, in that probability estimates are only available for quite a small number of variables (real GDP growth, core CPI inflation, core PCE inflation and the unemployment rate). Moreover, the horizons and levels of aggregation at which the participants are required to provide probability estimates often do not coincide with those for which they report point forecasts, something that is particularly evident in the case of real GDP. As such, we focus on the latter three variables, namely, core CPI inflation (CCPI), core PCE inflation (CPCE) and the unemployment rate (UNEMPR).

For CCPI and CPCE, participants report their perceived probabilities of the fourth-quarter-over-fourth-quarter changes in these variables falling into ten different ranges, both for the current year and the next year. For the purposes of this paper, we focus on the current-year 4Q-on-4Q forecast. These ranges have remained constant for the entire history of this survey question. For UNEMP, the same style of question is asked, but participants are tasked with reporting their probability estimates for four years (the current year and the next three years). The value ranges (bins), however, have changed over time, and three vintages of these bins exist in the data. It is also important to note that within a year, forecasters are tasked with producing *fixed-event* forecasts, in the sense that they are forecasting the same quantity over all four quarters of a year. Some previous work, such as Giordani and Söderlind (2003), has only focused on data from a single quarter in each year, but we find that the seasonal dynamics of fixed-event forecasts provide additional evidence to discriminate between different models of imperfect information.

Given that survey participants do not report a full continuous probability density forecast

(rather, we have discrete bins), it is important to come up with a consistent protocol to interpolate within bins. Here, we choose the *maximum-variance* interpolation, where we assume that each bin is drawn with the probability specified by the forecaster, and conditional on drawing a given bin, the variable is distributed uniformly. This is not only useful from the perspective of providing the most conservative representation of survey participants' uncertainty, but it also enables us to easily compute moments of the implied distribution, such as the mean and the variance. This approach is evidently not the only way to proceed, but is quite standard in the literature (see Zarnowitz and Lambros (1987) for a discussion). Indeed, with N bins and given bin probabilities $p_1, p_2, ..., p_N$, and left and right endpoints $L_1, R_1, L_2, R_2, ..., L_N, R_N$, the implied mean μ can be computed as $\sum_{i=1}^{N} \frac{1}{2}(R_i - L_i)p_i$ and the implied variance can be computed using the law of total variance. The conditional variance in each bin i is $V_i = \frac{1}{12}(R_i - L_i)^2$ and the conditional expectation in each bin is $E_i = \frac{1}{2}(R_i - L_i)$. We compute the implied variance (which corresponds to a forecaster's uncertainty) as $\mathrm{E}(V|.) + \mathrm{V}(E|.) = \sum_{i=1}^{N} V_i p_i + \sum_{i=1}^{N} E_i^2 p_i - \mu^2$. For each variable, we can then compute the average uncertainty across forecasters at time t, which we call U_t.

It is comparatively straightforward to compute cross-sectional disagreement, which we obtain by taking the variance of point predictions for a given variable across forecasters. Let us call this metric D_t. We now show some graphs and summary statistics that capture the dynamics and relationships between uncertainty and disagreement for the variables of interest.

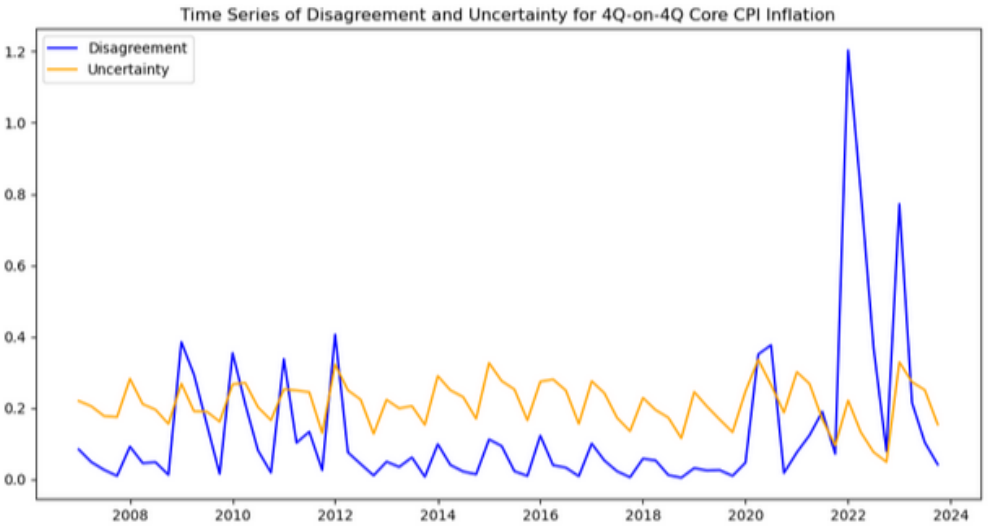


Figure 1: Dynamics of uncertainty and disagreement for core CPI inflation forecasts in the SPF

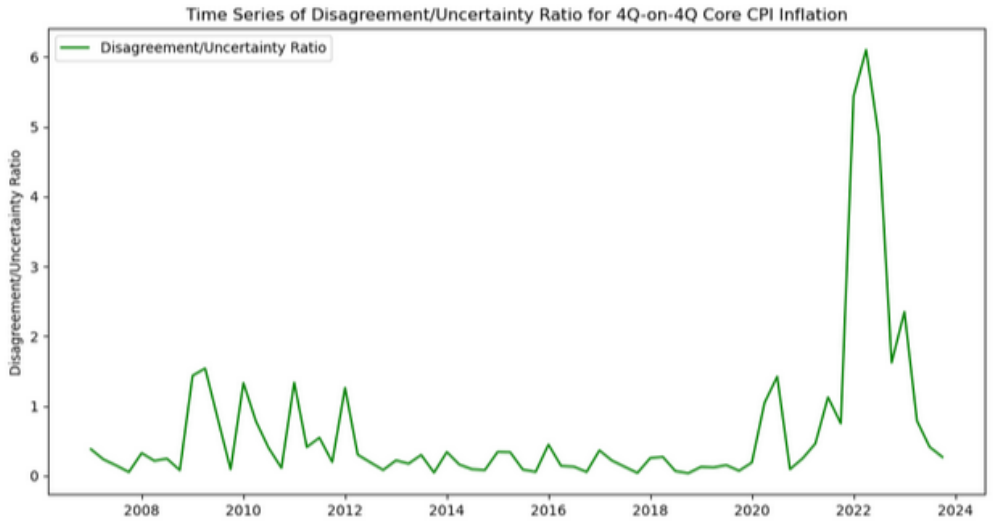


Figure 2: Dynamics of disagreement/uncertainty ratio for core CPI inflation forecasts in the SPF

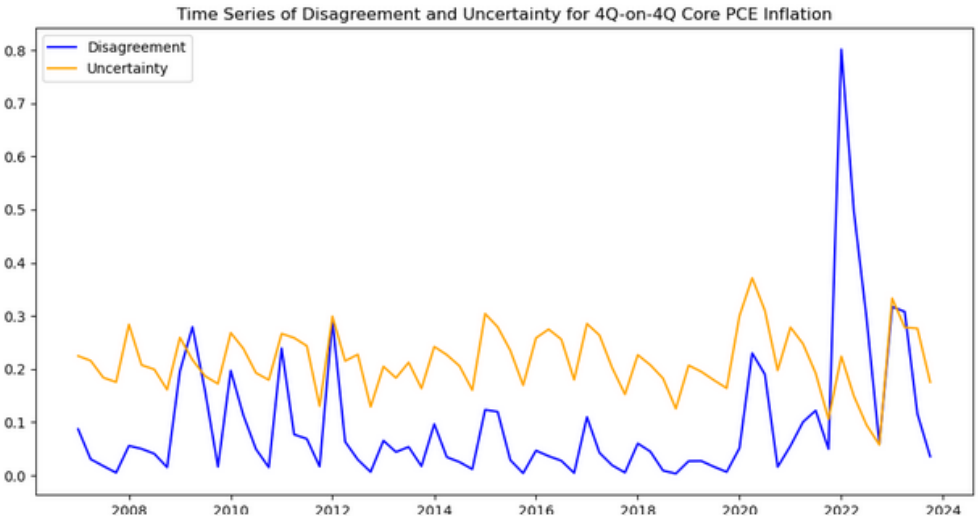


Figure 3: Dynamics of uncertainty and disagreement for core PCE inflation forecasts in the SPF

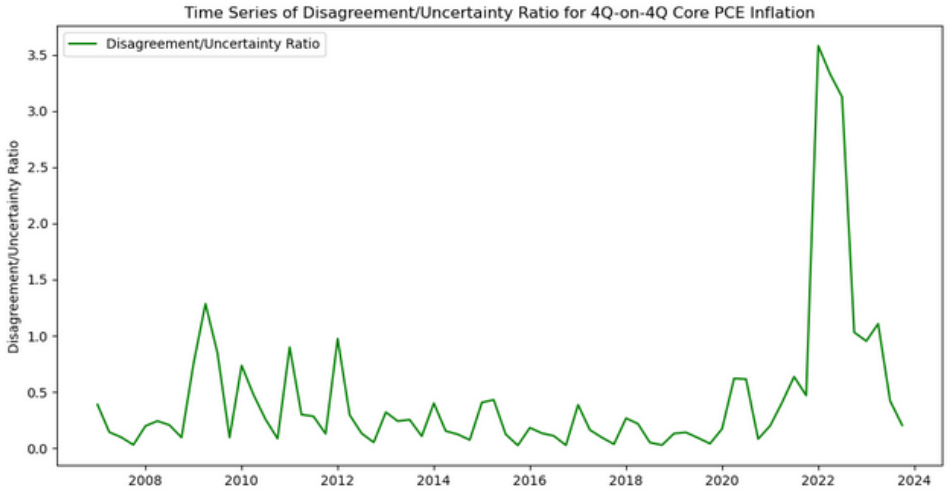


Figure 4: Dynamics of disagreement/uncertainty ratio for core PCE inflation forecasts in the SPF

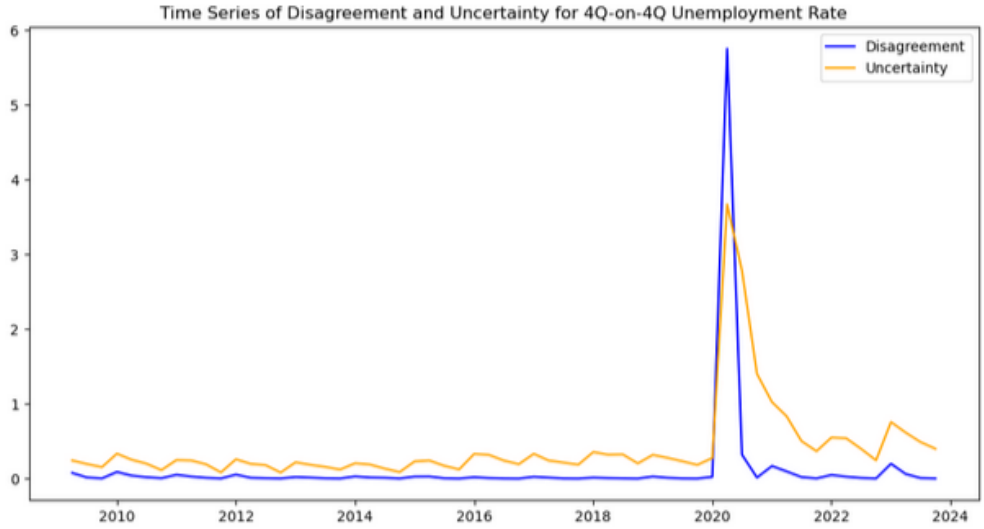


Figure 5: Dynamics of uncertainty and disagreement for unemployment forecasts in the SPF

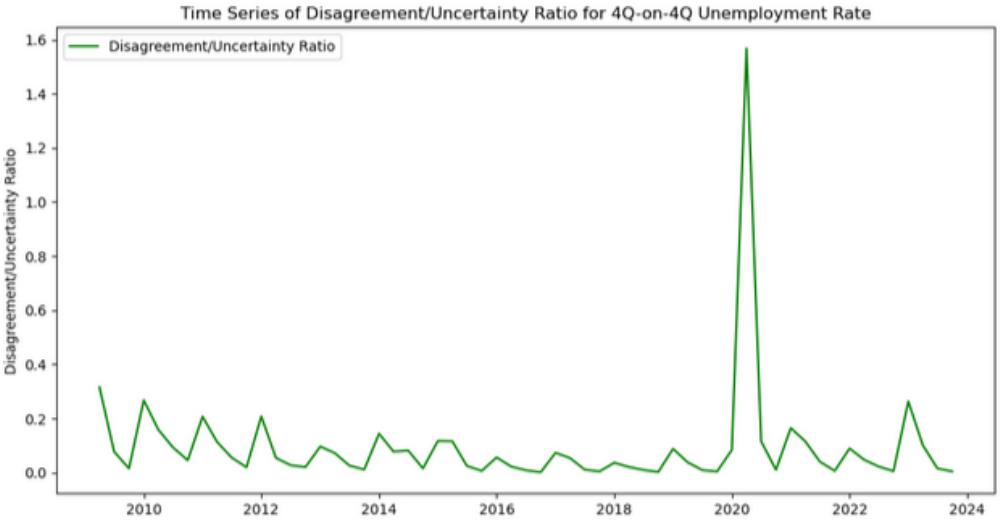


Figure 6: Dynamics of disagreement/uncertainty ratio for unemployment forecasts in the SPF

Barring the rather atypical COVID-19 period, some general observations from Figures 1 to 6 are that the disagreement/uncertainty ratio is generally less than unity, and even when it occasionally surpasses one (which typically appears to occur during periods of heightened macroeconomic volatility, such as the Great Recession or the COVID-19 pandemic), it does not do so by a large margin. Moreover, disagreement and uncertainty exhibit quite strong comovement (the pre-COVID correlation is 0.55 for core CPI inflation, 0.60 for core PCE

inflation, and 0.44 for unemployment), findings that are in line with those of previous authors, such as Zarnowitz and Lambros (1987) and Giordani and Söderlind (2003). There is also a striking seasonal pattern in uncertainty, disagreement and the uncertainty/disagreement ratio, which is natural given that the forecasters are tasked with making fixed-event forecasts over the four quarters of the year. Indeed, all of these three quantities fall between Q1 and Q4 (that is, disagreement falls much more in proportional terms than uncertainty). Motivated by these empirical observations, we will derive theoretical properties of uncertainty and disagreement and subsequently explore how well standard theoretical models of informational frictions in macroeconomics can rationalize these findings.

3 Theoretical Preliminaries

The simplest setting that captures the theoretical ideas in this paper is a static model, where the true state is given by the random variable X and the observed stochastic signal is given by the random variable Y. Consider an infinite continuum of forecasters who draw signals independently according to this probability distribution. A forecaster's uncertainty is simply the posterior variance of the state given the signal, that is, $\mathbb{V}(X|Y)$, and expected uncertainty is thus given by $\mathbb{E}[\mathbb{V}(X|Y)]$. Disagreement (which is synonymous with dispersion for the purposes of this paper), on the other hand, corresponds to the cross-sectional variance of forecasters' point forecasts conditional on the true state, that is, $\mathbb{V}[\mathbb{E}(X|Y)]$. Hence, expected disagreement is given by $\mathbb{E}[\mathbb{V}(\mathbb{E}(X|Y)|X)]$.

We begin to examine the relationship between these quantities by deriving a novel variance decomposition that captures the contributions of disagreement and uncertainty to the total variance of the true state variable. We first obtain from the law of total variance that

$$\mathbb{V}(X) = \underbrace{\mathbb{E}[\mathbb{V}(X|Y)]}_{\text{expected uncertainty}} + \mathbb{V}[\mathbb{E}(X|Y)]$$

The important point to note here is that the second term is *not* a suitable measure of disagreement, since the variance is unconditional. To see this, consider, for example, the scenario in which forecasters observe the true state noiselessly but there is variability in the true state. In this scenario, there would clearly be no disagreement across forecasters, but the second term in the variance decomposition would be positive. Instead, we will proceed by decomposing the second term with another application of the law of total variance, conditioning on the true state:

$$\mathbb{V}(X) = \underbrace{\mathbb{E}\left[\mathbb{V}(X|Y)\right]}_{\text{expected uncertainty}} + \underbrace{\mathbb{E}\left[\mathbb{V}(\mathbb{E}(X|Y)|X)\right]}_{\text{expected disagreement}} + \underbrace{\mathbb{V}\left[\mathbb{E}(\mathbb{E}(X|Y)|X)\right]}_{\text{forecast variability}}$$

In essence, the conditioning on the true state sequence is necessary to avoid misattributing unconditional movement in the forecasts as disagreement. While this simple setup encompasses many standard models, for example, the one-period noisy and sticky information models in macroeconomics, it is not well suited to handle more complex signal structures. For example, forecasters might form expectations based on the full history of signals observed thus far (as in the dynamic noisy information model), or might be stochastically stuck with stale information (as in the dynamic sticky information model). To handle a far more general class of informational structures, we proceed as follows.

Let $X \in L^2$ be a random variable defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and suppose $\mathcal{G}$ and $\mathcal{H}$ are σ-algebras such that $\mathcal{G}, \mathcal{H} \subset \mathcal{F}$ and $\sigma(X) \subset \mathcal{G}$. Here, $\mathcal{G}$ represents shared/predetermined information across forecasters (even though they may not know this information), while $\mathcal{H}$ represents forecasters' private information. We make the natural assumption that $\sigma(X) \subset \mathcal{G}$ (in essence, the true state is "stochastically predetermined" prior to the generation of private information sets).

In the subsequent sections, we now proceed to analyze the relative magnitudes of uncertainty and disagreement under different informational regimes. In particular, in light of the empirical evidence, we explore whether it is possible for disagreement to be greater than

uncertainty, as well as what conditions lead to temporal and state-dependent variation in these quantities.

4 Main Theorem

In this section, we present the key result of this paper, a theorem that states that expected disagreement is always less than or equal to expected uncertainty when forecasters have a common prior and form rational forecasts based on private signals.

Theorem 1. *Given an L^2-integrable random variable X and σ-algebras $\mathcal{G}$ and $\mathcal{H}$ on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$, if $\sigma(X) \subset \mathcal{G}$, we have that $\mathbb{E}[\mathbb{V}[\mathbb{E}(X|\mathcal{H})|\mathcal{G}]] \leq \mathbb{E}[\mathbb{V}(X|\mathcal{H})]$.*

This theorem provides a rigorous formulation of the previously discussed concepts of uncertainty and disagreement using the language of σ-algebras. To prove the result, we begin by proving an auxiliary lemma.

Lemma 1. *For any L^2-integrable random variable X and σ-algebra $\mathcal{G}$ on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$, we have that $\mathbb{V}[X - \mathbb{E}(X|\mathcal{G})] = \mathbb{E}[\mathbb{V}(X|\mathcal{G})]$. Furthermore, $\mathbb{V}[X - \mathbb{E}(X|\mathcal{G})] \leq \mathbb{V}(X)$.*

Proof. Let $U = X - \mathbb{E}(X|\mathcal{G})$. For the first part of the result, we proceed as follows:

$$\begin{aligned}
\mathbb{V}(U) &= \mathbb{E}[\mathbb{V}(U|\mathcal{G})] + \mathbb{V}[\mathbb{E}(U|\mathcal{G})] \text{ (by the law of total variance)} \\
&= \mathbb{E}[V(X|\mathcal{G})] + 0 \\
&= \mathbb{E}[V(X|\mathcal{G})]
\end{aligned}$$

Now, for the second part, we have that:

$$\begin{aligned}
\mathbb{V}(X) &= \mathbb{E}[\mathbb{V}(X|\mathcal{G})] + \mathbb{V}[\mathbb{E}(X|\mathcal{G})] \text{ (by the law of total variance)} \\
&= \mathbb{V}(U) + \mathbb{V}[\mathbb{E}(X|\mathcal{G})] \text{ (by the first part of the lemma)}
\end{aligned}$$

$$\geq \mathbb{V}(U)$$
$$= \mathbb{V}[X - \mathbb{E}(X|\mathcal{G})]$$

□

We now proceed to the proof of the main theorem:

Proof.

$$\begin{aligned}
\mathbb{E}[\mathbb{V}[\mathbb{E}(X|\mathcal{H})|\mathcal{G}]] &= \mathbb{V}\left[\mathbb{E}(X|\mathcal{H}) - \mathbb{E}[\mathbb{E}(X|\mathcal{H})|\mathcal{G}]\right] \text{ (by the first part of Lemma 1)} \\
&= \mathbb{V}\left[X - \mathbb{E}(X|\mathcal{H}) - \mathbb{E}[X - \mathbb{E}(X|\mathcal{H})|\mathcal{G}]\right] \text{ (since } \sigma(X) \subset \mathcal{G}) \\
&\leq \mathbb{V}[X - \mathbb{E}(X|\mathcal{H})] \text{ (by the second part of Lemma 1)} \\
&= \mathbb{E}[\mathbb{V}(X|\mathcal{H})] \text{ (by another application of the first part of Lemma 1)}
\end{aligned}$$

□

A useful (but admittedly incomplete) visual representation of the ideas in the proof is shown below in Figure 7.

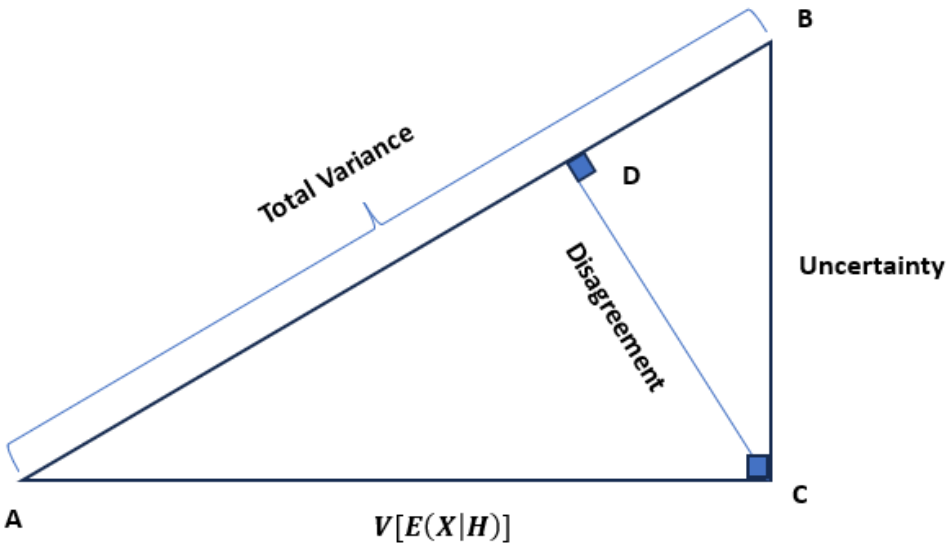


Figure 7: Geometric visualization of proof of main theorem

In essence, line segment CD is the common orthogonal rejection of $\mathbb{E}(X|\mathcal{H})$ and $X - \mathbb{E}(X|\mathcal{H})$ on $\mathcal{G}$, which enables us to prove the desired result. However, it is not necessarily true in the general case that the projections on $\mathcal{G}$ are collinear, as suggested by the figure (this appears to be only true when we restrict attention to univariate linear projections). Nonetheless, the visualization offers much intuition into the nature of the proof strategy.

Indeed, a further implication of the figure is that triangles BDC and BCA are similar, which implies that $\frac{D}{U} = \frac{V-U}{V} = 1 - \frac{U}{V}$, where D represents expected disagreement, U represents expected uncertainty and V represents the total variance of X. In other words, one might conjecture that the ratio of expected disagreement to expected uncertainty is decreasing in the degree of informational frictions that forecasters face. While this is of course far from a rigorous argument, it is interesting that this underlying intuition indeed turns out to hold in the standard macroeconomic models of noisy information and sticky information that we explore later in the paper.

It is also valuable to consider which assumptions can and cannot be relaxed while still preserving the key result. For example, suppose that there are some public signals in addition to the existing private signals that forecasters observe. Then we replace the statement of Theorem 1 by $\mathbb{E}[\mathbb{V}[\mathbb{E}(X|\mathcal{G}^{\mathcal{C}})|\mathcal{F}^{\mathcal{C}}]] \leq \mathbb{E}[\mathbb{V}(X|\mathcal{G}^{\mathcal{C}})]$, where the superscript $\mathcal{C}$ denotes a refinement of the original σ-algebra with the additional information encoded by the public signals. Since $\sigma(X) \subset \mathcal{G} \subset \mathcal{G}^{\mathcal{C}}$, the theorem still holds in this case. Moreover, the theorem allows for heterogeneity in forecasters' signal distributions, since we are effectively considering the forecasters as an infinite population and taking cross-sectional averages in computing disagreement and average uncertainty.

Moreover, the theorem also accounts for settings where there is *irreducible uncertainty*, that is, when even the union of forecasters' information sets is insufficient to identify the true state. For example, suppose forecasters obtained noisy signals about a latent state, where the true state is this latent state plus some independent noise. The σ-algebra $\mathcal{G}$ could encode both the latent and true states, preserving the validity of the theorem.

However, it is easy to see that the theorem need not hold if expectations are non-rational. If, say, forecasters always reported their noisy signal as their best guess, where this signal is the true state plus some random noise, then expected disagreement could certainly exceed expected uncertainty. In some sense, the problem is not even well-posed, since we would need to define how the posterior uncertainty of forecasters should be computed.

An additional point to note is that the theorem only considers disagreement and uncertainty in expectation. It does not say that disagreement is always less than uncertainty for any given state of the world. Indeed, we will see that under the model of sticky information, disagreement can exceed uncertainty when large shocks occur. Nonetheless, the theorem is suggestive that disagreement exceeding uncertainty is an atypical outcome, an intuition that we will formalize for some specific settings later in the paper.

However, it is possible to derive a conditional (or "pointwise") version of the theorem based on realized forecasts and forecast errors. If we let Z denote the random forecast (which could be rational or irrational), we can derive that:

$$\underbrace{\mathbb{E}[(X-Z)^2|\mathcal{G}]}_{\text{expected realized error}} = \underbrace{\mathbb{E}[(X-\mathbb{E}(Z|\mathcal{G}))^2|\mathcal{G}]}_{\text{consensus error}} + \underbrace{\mathbb{E}[(\mathbb{E}(Z|\mathcal{G})-Z)^2|\mathcal{G}]}_{\text{disagreement}}$$

In this result, the realized error always exceeds disagreement for any state of the world, which implies that if uncertainty is conditionally less than disagreement in a given state, then it must be an underestimate of the expected realized forecast error.

Naturally, to fully characterize the dynamics of disagreement and uncertainty, we require parametric assumptions on the stochastic processes governing the true state and the signals observed by forecasters. As such, we now proceed to investigate special cases of this theorem with macroeconomic applications, namely, variants of Woodford (2003)'s model of noisy information and Mankiw and Reis (2002)'s model of sticky information. With the specific assumptions of these models, we can derive other complementary comparative statics results and ascertain a broad suite of theoretical implications of these models for the dynamics of

disagreement and uncertainty. We will then explore the relative ability of these two models to fit the empirical behavior of disagreement and uncertainty in the SPF data.

5 Macroeconomic Applications: Static Models

In this section, we analyze simple versions of standard macroeconomic models with informational frictions where there is no persistence in the state variable. As such, in each period t, the true state $X_t = \mu_X + \epsilon_t$, where $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ and ϵ_t is independent and identically distributed over time. Without loss of generality, we may normalize $\mu_X = 0$. Under the first informational regime that we study, namely, *noisy information*, each forecaster observes a noisy signal of the true state, $Y_{it} = X_t + \nu_{it}$, where $\nu_{it} \sim N(0, \sigma_\nu^2)$ and is independent across forecasters and over time. Since we consider forecasters as an "infinite population", we may proceed to drop the i subscript, with the understanding that the realization is for an arbitrary forecaster.

We are interested in deriving expressions for uncertainty and disagreement, as defined in the previous section, in this environment. Applying standard results pertaining to the multivariate normal distribution and noting that forecasters need only condition on Y_t since there is no persistence in the state, we find that the conditional variance of X_t given Y_t, which we define by U_t (since it denotes forecast uncertainty) is $(1-\rho^2)\sigma_\epsilon^2$, where $\rho = Corr(X_t, Y_t) = \frac{\sigma_\epsilon^2}{\sigma_\epsilon\sqrt{\sigma_\epsilon^2+\sigma_\nu^2}} = \frac{\sigma_\epsilon}{\sqrt{\sigma_\epsilon^2+\sigma_\nu^2}}$. Hence, uncertainty, which is the conditional variance, is given by

$$U_t = \frac{\sigma_\epsilon^2 \sigma_\nu^2}{\sigma_\epsilon^2 + \sigma_\nu^2}$$

. Note that if we fix the underlying state variance σ_ϵ^2, we find that as the noise variance σ_ν^2 tends to 0, the conditional variance also tends to 0. Moreover, uncertainty is increasing in σ_ν^2 and tends to σ_ϵ^2 as σ_ν^2 tends to ∞. The symmetric argument applies if we fix σ_ν^2 and vary σ_ϵ^2, in that uncertainty is increasing in σ_ϵ^2 and tends to σ_ν^2 as σ_ϵ^2 tends to ∞ . Another interesting but intuitive finding is that uncertainty is upper bounded by the maximum of

the signal variance and the noise variance.

Now, let us examine disagreement across forecasters in their point forecasts, which, following the previous section, is defined as the variance of the posterior mean conditional on the true sequence of states $\{X_t\}$. Again, because of the absence of any persistence in the state, we need only condition on the current state X_t. Again using standard properties of the multivariate normal distribution, we find that the posterior mean is $\frac{\rho\sigma_\epsilon}{\sqrt{\sigma_\epsilon^2+\sigma_\nu^2}}Y_t = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\nu^2}Y_t = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\nu^2}(X_t+\nu_t)$. Conditional on X_t, the stochastic component is $\frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\nu^2}\nu_t$. Thus, the cross-sectional forecast disagreement is given by

$$D_t = \frac{\sigma_\epsilon^4\sigma_\nu^2}{(\sigma_\epsilon^2+\sigma_\nu^2)^2}$$

. Moreover, we have that $(\sigma_\epsilon^2-\sigma_\nu^2)^2 = (\sigma_\epsilon^2+\sigma_\nu^2)^2 - 4\sigma_\epsilon^2\sigma_\nu^2 \geq 0$, so that $\frac{\sigma_\epsilon^2\sigma_\nu^2}{(\sigma_\epsilon^2+\sigma_\nu^2)^2} \leq \frac{1}{4}$. Hence, for a fixed σ_ϵ^2, we have that $D_t \leq \frac{\sigma_\epsilon^2}{4}$, and this bound is attained with equality when $\sigma_\nu^2 = \sigma_\epsilon^2$. Thus, disagreement is maximized when the signal and noise are equally variable. This formalizes the intuition discussed in the introduction, where we highlighted that disagreement can be low both when individual signals are very precise (thereby essentially revealing the true state) as well as when individual signals are extremely noisy (whereupon forecasters essentially ignore their private signals and pool at the prior distribution).

A few useful observations can be made at this point. First, note that the disagreement-uncertainty ratio is given by

$$\frac{D_t}{U_t} = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\nu^2},$$

implying that disagreement cannot exceed uncertainty, which accords with the finding of the general theorem. Equality holds if and only if we have the degenerate case of $\sigma_\nu^2 = 0$, that is, when the signals observed by the forecasters are noiseless. Second, neither uncertainty nor disagreement are time-varying or dependent on the realizations of either the systematic shock ϵ_t or the idiosyncratic noise shock ν_t.

Now, let us consider the second canonical model of information frictions discussed in the

macroeconomics literature, namely, *sticky information*. Under this informational regime, forecasters' observations of the true state are not contaminated with noise, but instead, with probability $1 - \lambda$, they update their information set at each time point t and with complementary probability λ, their information remains stale. Thus, λ serves as an index of informational rigidities.

Let us begin by examining uncertainty in this sticky information environment. Observe that all forecasters who updated their information set in the current period have zero uncertainty (a mass of $1-\lambda$), while the uncertainty of forecasters with stale information (a mass of λ) is given by σ_ϵ^2. Immediately, we see that unlike the setting with noisy information, there is cross-sectional heterogeneity in forecasters' uncertainty. Averaging across forecasters, the mean uncertainty is given by $U_t = \lambda\sigma_\epsilon^2$. Observe that this quantity is not time-varying, since we are considering the forecasters at a population level.

Next, let us look at inter-forecaster disagreement, that is, the dispersion of point forecasts across forecasters. Indeed, the point forecast of all forecasters who update their information set in the current period (a mass of $1-\lambda$) is ϵ_t (recall that we normalized $\mu_X = 0$), while the point forecast of all other forecasters (the mass λ with stale information) is 0. Hence, the variance of the point forecast is $D_t = \lambda(1-\lambda)\epsilon_t^2$. Thus, we see that cross-sectional forecast disagreement is time-varying and increasing in the magnitude of the shock under sticky information. Moreover, the unconditional expectation of disagreement is $D_u = \lambda(1-\lambda)\sigma_\epsilon^2$, where the u subscript denotes that we are looking at the unconditional version of this metric.

Therefore, we find that $\frac{D_u}{U_t} = 1 - \lambda$, so it emerges that in an unconditional sense, disagreement again cannot be greater than uncertainty under sticky information in this static environment, which is again an illustration of the main theorem. Uncertainty and disagreement are equal if and only if $\lambda = 0$, that is, when forecasters have perfect information. Thus, there are qualitatively similar relationships between uncertainty and disagreement in the cases of noisy and sticky information, but the underlying mechanisms are, of course, very different. Under both noisy information and sticky information, uncertainty and dis-

agreement are closest in magnitude when information frictions are at their lowest, that is, when observations are noiseless or when there is no informational stickiness. In both cases, the disagreement-uncertainty ratio is decreasing in the degree of informational frictions (as we predicted from the non-rigorous geometric argument in the previous section). However, it is worth emphasizing an important difference between the sticky and noisy information environments. Due to state-dependent variation in cross-sectional forecast disagreement, it is possible under sticky information for disagreement to exceed uncertainty at a given point in time. Indeed, the ratio of conditional disagreement to uncertainty is $\frac{(1-\lambda)\epsilon_t^2}{\sigma_\epsilon^2}$. Therefore, disagreement exceeds uncertainty if and only $\epsilon_t^2 \geq \frac{\sigma_\epsilon^2}{1-\lambda}$.

6 Macroeconomic Applications: Dynamic Models

We now consider a more complex dynamic setting where the underlying state follows a stationary AR(1) process. Therefore, we have that $X_t = \phi X_{t-1} + \epsilon_t$, where $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ is independent and identically distributed over time and $0 \leq \phi < 1$. As in the previous section, we begin by considering the case of noisy information, where each forecaster observes the noisy signal $Y_{it} = X_t + \nu_{it}$ and $\nu_{it} \sim N(0, \sigma_\nu^2)$ is independent across forecasters and over time. Again, since we consider the forecasters as an infinite population, we will proceed to suppress the i subscript.

In this environment, forecasters apply the Kalman filter to generate posterior distributions for the current state given the observed data. At each time step t, the Kalman filter algorithm performs the following steps:

$$\hat{X}_{t|t-1} = \phi \hat{X}_{t-1|t-1}$$

$$P_{t|t-1} = \phi^2 P_{t-1} + \sigma_\epsilon^2$$

$$K_t = \frac{P_{t|t-1}}{P_{t|t-1} + \sigma_\nu^2}$$

$$\hat{X}_{t|t} = \hat{X}_{t|t-1} + K_t(Y_t - \hat{X}_{t|t-1})$$

$$P_t = (1 - K_t)P_{t|t-1}$$

These equations are simply the standard Kalman filter equations adapted to our univariate setting. X_t is the posterior mean of the state at time t given the information set at time t, $X_{t|t-1}$ is the posterior mean of the state at time t given the information set at time $t-1$, P_t is the posterior variance of the state at time t given the information set at time t, $P_{t|t-1}$ is the posterior variance of the state at time t given the information set at time $t-1$, and K_t is the *Kalman gain*, which specifies the relative weight that should be put on the new information at time t.

From Hamilton (1994), we know that with our given assumptions, the Kalman filter converges to a steady state, which we find by solving the *Riccati equation* $P_{t-1} = P_t = P_\infty$, replacing K_t with the steady-state Kalman gain K_∞. Hence, $K_\infty = \frac{\phi^2 P_\infty + \sigma_\epsilon^2}{\phi^2 P_\infty + \sigma_\epsilon^2 + \sigma_\nu^2}$ and $P_\infty = \frac{\sigma_\nu^2(\phi^2 P_\infty + \sigma_\epsilon^2)}{\phi^2 P_\infty + \sigma_\epsilon^2 + \sigma_\nu^2}$. The last equality yields the following quadratic equation in P_∞:

$$\phi^2 P_\infty^2 + \left[\sigma_\epsilon^2 + \sigma_\nu^2(1-\phi^2)\right] P_\infty - \sigma_\epsilon^2 \sigma_\nu^2 = 0$$

Clearly, we must take the positive root, so that

$$P_\infty = \frac{-\sigma_\epsilon^2 - \sigma_\nu^2(1-\phi^2) + \sqrt{\left[\sigma_\epsilon^2 + \sigma_\nu^2(1-\phi^2)\right]^2 + 4\phi^2\sigma_\epsilon^2\sigma_\nu^2}}{2\phi^2}$$

It is precisely P_∞ that is the steady-state uncertainty of a given forecaster, and once again, it is time and state-invariant.

Next, we look at the cross-sectional disagreement of point forecasts under the noisy information regime in this dynamic environment. Analogous to the simpler static environment, it is natural to define this as the steady-state variance of the posterior distribution conditional on the true sequence of states $\{X_t\}$. However, the dynamic case is somewhat more complicated as forecasters need to genuinely condition on the full sequence of states.

To proceed, consider the update rule of the Kalman filter in steady state:

$$\hat{X}_{t|t} = \phi\hat{X}_{t-1|t-1} + K_\infty(Y_t - \phi\hat{X}_{t-1|t-1}) = \phi(1-K_\infty)\hat{X}_{t-1|t-1} + K_\infty(X_t + \nu_t)$$

From the convergence of the Kalman filter, we have that $\mathbb{V}(\hat{X}_{t|t}|\{X_t\})$ converges in the ergodic regime to $\mathbb{V}_\infty$, which we find by solving $\mathbb{V}(\hat{X}_{t|t}|\{X_t\}) = \mathbb{V}(\hat{X}_{t-1|t-1}|\{X_t\}) = \mathbb{V}(\hat{X}_{t-1|t-1}|\{X_{t-1}\}) = \mathbb{V}_\infty$, we can substitute this into the update rule and take the conditional variance of both sides with respect to $\{X_t\}$ (noting that X_t is in the conditioning set) to obtain $\mathbb{V}_\infty = \phi^2(1-K_\infty)^2\mathbb{V}_\infty + K_\infty^2\sigma_\nu^2$. Therefore, we have that

$$\mathbb{V}_\infty = \frac{K_\infty^2\sigma_\nu^2}{1-\phi^2(1-K_\infty)^2}$$

Since $\mathbb{V}_\infty$ corresponds to disagreement and P_∞ corresponds to uncertainty, it is convenient to instead define $D_\infty \equiv \mathbb{V}_\infty$ and $U_\infty \equiv P_\infty$. Note that in the static case, we have $\phi = 0$ and $K_\infty = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\nu^2}$, so that we indeed recover that $D_\infty = \frac{\sigma_\epsilon^4\sigma_\nu^2}{(\sigma_\epsilon^2+\sigma_\nu^2)^2}$

We now proceed to plot $R \equiv \frac{D_\infty}{U_\infty}$ for different values of ϕ, σ_ϵ^2 and σ_ϵ^2. As Figures 8 and 9 show, we see that $\frac{D_\infty}{U_\infty}$ appears to be increasing in ϕ, decreasing in σ_ν^2 and increasing in σ_ϵ^2. Moreover, the disagreement-uncertainty ratio appears to tend to 1 as σ_ϵ^2 becomes large, as well as when σ_ν^2 tends to 0, but it seems that disagreement never exceeds uncertainty. As it turns out, it is in fact possible to analytically prove all of these properties.

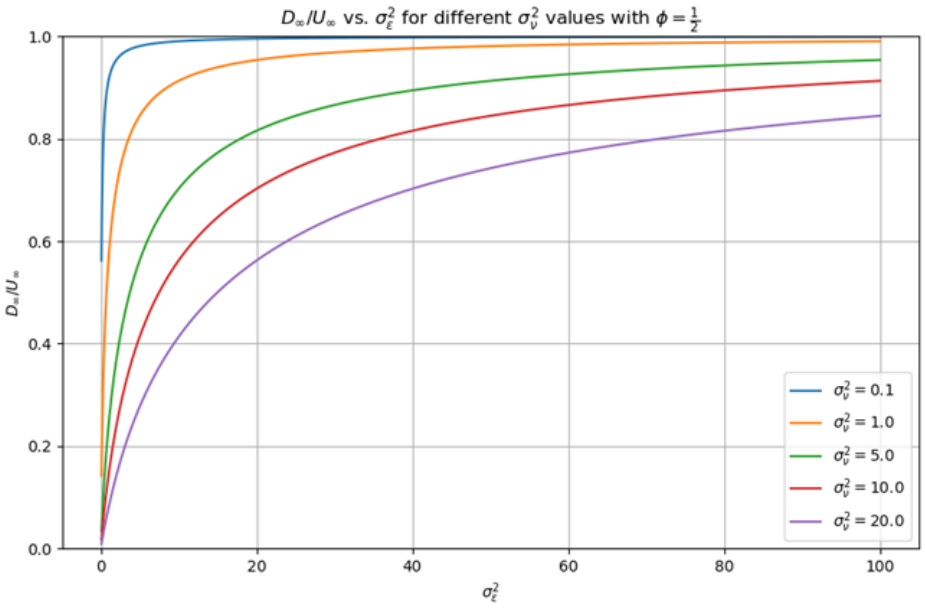


Figure 8: Comparative statics for disagreement/uncertainty ratio with respect to σ_ϵ^2 and σ_ν^2 under noisy information

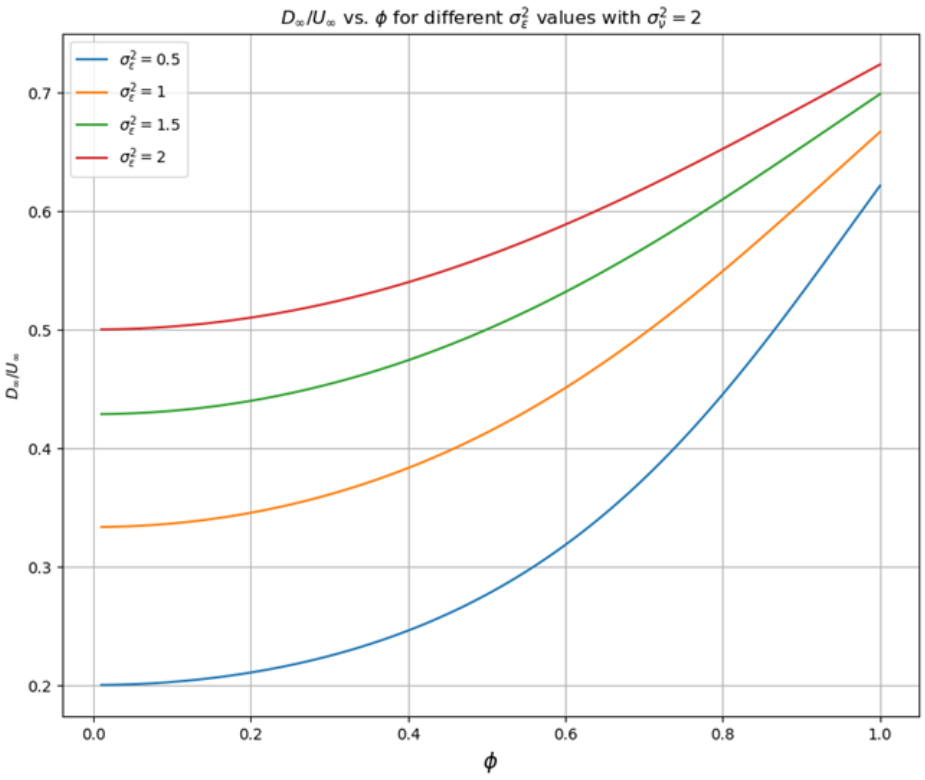


Figure 9: Comparative statics for disagreement/uncertainty ratio with respect to ϕ under noisy information

Proposition 1. *With noisy information, $R \equiv \frac{D_\infty}{U_\infty}$ is increasing in ϕ, decreasing in σ_ν^2 and increasing in σ_ϵ^2. Moreover, $0 \leq R \leq 1$ and $\lim_{\sigma_\epsilon^2 \to \infty} R = \lim_{\sigma_\nu^2 \to 0} R = 1$.*

Proof. After some algebraic manipulation, we find that:

$$\frac{dR}{d\sigma_\epsilon^2} = \frac{\sigma_\nu^2 \left(\sigma_\epsilon^2 + \sigma_\nu^2(1-\phi^2)\right)}{\left(4\phi^2\sigma_\epsilon^2\sigma_\nu^2 + \left(\sigma_\epsilon^2 + \sigma_\nu^2 - \phi^2\sigma_\nu^2\right)^2\right)^{3/2}} > 0$$

$$\frac{dR}{d\sigma_\nu^2} = -\frac{\sigma_\epsilon^2 \left(\sigma_\epsilon^2 + \sigma_\nu^2 - \phi^2\sigma_\nu^2\right)}{\left(\sigma_\epsilon^4 + 2(1+\phi^2)\sigma_\epsilon^2\sigma_\nu^2 + (-1+\phi^2)^2\sigma_\nu^4\right)^{3/2}} < 0$$

$$\frac{dR}{d\phi} = -\frac{4\phi\sigma_\epsilon^2\sigma_\nu^4}{\left(4\phi^2\sigma_\epsilon^2\sigma_\nu^2 + \left(\sigma_\epsilon^2 + \sigma_\nu^2 - \phi^2\sigma_\nu^2\right)^2\right)^{3/2}} > 0$$

Thus, the first three properties are directly proven. To show the asymptotic behavior, first observe that $\lim_{\sigma_\epsilon^2 \to \infty} K_\infty = 1$, so that $\lim_{\sigma_\epsilon^2 \to \infty} D_\infty = \sigma_\nu^2$. All that remains is to characterize the asymptotic behavior of U_∞. Indeed, we have:

$$\begin{aligned}
U_\infty &= \frac{-\sigma_\epsilon^2 - \sigma_\nu^2(1-\phi^2) + \sqrt{\left[\sigma_\epsilon^2 + \sigma_\nu^2(1-\phi^2)\right]^2 + 4\phi^2\sigma_\epsilon^2\sigma_\nu^2}}{2\phi^2} \\
&= \frac{-\sigma_\epsilon^2 - \sigma_\nu^2(1-\phi^2) + \sqrt{\sigma_\epsilon^4 + \sigma_\nu^4(1-\phi^2)^2 + 2\sigma_\epsilon^2\sigma_\nu^2(1-\phi^2) + 4\phi^2\sigma_\epsilon^2\sigma_\nu^2}}{2\phi^2} \\
&\sim \frac{-\sigma_\epsilon^2 - \sigma_\nu^2(1-\phi^2) + \sqrt{\sigma_\epsilon^4 + 2\sigma_\epsilon^2\sigma_\nu^2(1-\phi^2) + 4\phi^2\sigma_\epsilon^2\sigma_\nu^2}}{2\phi^2} \\
&= \frac{-\sigma_\epsilon^2 - \sigma_\nu^2(1-\phi^2) + \sqrt{\sigma_\epsilon^4 + 2\sigma_\epsilon^2\sigma_\nu^2(1+\phi^2)}}{2\phi^2} \\
&= \frac{-\sigma_\epsilon^2 - \sigma_\nu^2(1-\phi^2) + \sigma_\epsilon^2\sqrt{1 + 2\frac{\sigma_\nu^2}{\sigma_\epsilon^2}(1+\phi^2)}}{2\phi^2} \\
&\sim \frac{-\sigma_\epsilon^2 - \sigma_\nu^2(1-\phi^2) + \sigma_\epsilon^2\left(1 + \frac{\sigma_\nu^2}{\sigma_\epsilon^2}(1+\phi^2)\right)}{2\phi^2} \\
&= \sigma_\nu^2
\end{aligned}$$

Here, $\sim$ denotes asymptotic equivalence, and the second-last line follows from a first-order Taylor series expansion of the function $f(x) = \sqrt{1+x}$. Using this result, we find that indeed $\frac{D_\infty}{U_\infty}$ tends to 1 as σ_ϵ^2 tends to ∞. Moreover, monotonicity with respect to σ_ϵ^2 and this

limiting behavior imply that D_∞ never exceeds U_∞. Finally, the same asymptotic behavior for $P = D_\infty$ and U_∞ occurs as σ_ν^2 tends to 0, so we also have that $\frac{D_\infty}{U_\infty}$ tends to 1 as σ_ν^2 tends to 0, as required. □

Now, let us examine the case of sticky information. The point forecast of an agent who last updated their information set s periods ago is $\phi^s X_{t-s}$, and the concomitant steady-state mass of agents is $\lambda^s(1-\lambda)$.

Therefore, the conditional steady-state disagreement given the realized sequence of states is:

$$D_t = \sum_{s=0}^{\infty} \lambda^s(1-\lambda)\phi^{2s} X_{t-s}^2 - \left[\sum_{s=0}^{\infty} \lambda^s(1-\lambda)\phi^s X_{t-s}\right]^2$$

This is admittedly a complicated and time/state-dependent quantity that, unlike in the static case, exhibits dependence on the full sequence of previously realized states. Now, we will proceed to take the unconditional expectation to find the unconditional steady-state cross-sectional disagreement D_∞. Note first from the moments of the AR(1) process that $\mathbb{E}(X_t) = 0$, $\mathbb{V}(X_t) = \mathbb{E}(X_t^2) = \frac{\sigma_\epsilon^2}{1-\phi^2}$ and $\text{Cov}(X_{t-s}, X_{t-r}) = \mathbb{E}(X_{t-s}X_{t-r}) = \frac{\phi^{|s-r|}\sigma_\epsilon^2}{1-\phi^2}$. Now, the first term of D_∞ simplifies to:

$$\sum_{s=0}^{\infty} \lambda^s(1-\lambda)\phi^{2s}\frac{\sigma_\epsilon^2}{1-\phi^2} = \frac{(1-\lambda)\sigma_\epsilon^2}{(1-\lambda\phi^2)(1-\phi^2)}$$

The second term of D_∞ simplifies to:

$$\begin{aligned}
\frac{\sigma_\epsilon^2(1-\lambda)^2}{1-\phi^2}\sum_{r=0}^{\infty}\sum_{s=0}^{\infty}\lambda^{s+r}\phi^{s+r}\phi^{|s-r|} &= \frac{\sigma_\epsilon^2(1-\lambda)^2}{1-\phi^2}\left(2\sum_{r=0}^{\infty}\sum_{s=r+1}^{\infty}\lambda^{r+s}\phi^{2s} + \sum_{s=0}^{\infty}\lambda^{2s}\phi^{2s}\right) \\
&= \frac{\sigma_\epsilon^2(1-\lambda)^2}{1-\phi^2}\left(2\sum_{r=0}^{\infty}\frac{(\lambda\phi^2)^{r+1}\lambda^r}{1-\lambda\phi^2} + \frac{1}{1-\lambda^2\phi^2}\right) \\
&= \frac{\sigma_\epsilon^2(1-\lambda)^2}{1-\phi^2}\left(2\lambda\phi^2\sum_{r=0}^{\infty}\frac{(\lambda\phi)^{2r}}{1-\lambda\phi^2} + \frac{1}{1-\lambda^2\phi^2}\right) \\
&= \frac{\sigma_\epsilon^2(1-\lambda)^2}{1-\phi^2}\left(\frac{2\lambda\phi^2}{(1-\lambda\phi^2)(1-\lambda^2\phi^2)} + \frac{1}{1-\lambda^2\phi^2}\right)
\end{aligned}$$

$$= \frac{\sigma_\epsilon^2(1-\lambda)^2(1+\lambda\phi^2)}{(1-\phi^2)(1-\lambda\phi^2)(1-\lambda^2\phi^2)}$$

Hence,

$$\begin{aligned} D_\infty &= \frac{(1-\lambda)\sigma_\epsilon^2}{(1-\lambda\phi^2)(1-\phi^2)}\left(1 - \frac{(1-\lambda)(1+\lambda\phi^2)}{1-\lambda^2\phi^2}\right) \\ &= \frac{(1-\lambda)\sigma_\epsilon^2\lambda(1-\phi^2)}{(1-\lambda\phi^2)(1-\phi^2)(1-\lambda^2\phi^2)} \\ &= \frac{\lambda(1-\lambda)\sigma_\epsilon^2}{(1-\lambda\phi^2)(1-\lambda^2\phi^2)} \end{aligned}$$

Moreover, the individual uncertainty of a forecaster who last updated their information s periods ago is $\sum_{i=0}^{s-1}\phi^{2i}\sigma_\epsilon^2 = \sigma_\epsilon^2\frac{1-\phi^{2s}}{1-\phi^2}$, and as we noted earlier, the associated mass of forecasters is $\lambda^s(1-\lambda)$. Hence, the average uncertainty across forecasters is:

$$\begin{aligned} U_\infty &= (1-\lambda)\sigma_\epsilon^2\sum_{s=0}^{\infty}\lambda^s\frac{1-\phi^{2s}}{1-\phi^2} \\ &= (1-\lambda)\sigma_\epsilon^2\left[\frac{1}{(1-\lambda)(1-\phi^2)} - \frac{1}{(1-\lambda\phi^2)(1-\phi^2)}\right] \end{aligned}$$

Simplifying the latter expression, we have

$$U_\infty = \frac{\lambda\sigma_\epsilon^2}{(1-\lambda\phi^2)}$$

Hence,

$$R \equiv \frac{D_\infty}{U_\infty} = \frac{1-\lambda}{1-\lambda^2\phi^2}$$

The following proposition is now a straightforward corollary since $\frac{dR}{d\lambda} = \frac{-1+\phi^2(2\lambda-\lambda^2)}{(1-\lambda^2\phi^2)^2} < 0$:

Proposition 2. *With sticky information, $R \equiv \frac{D_\infty}{U_\infty}$ is increasing in ϕ and decreasing in λ. Moreover, $0 < R < 1$.*

Once again, we can plot $\frac{D_\infty}{U_\infty}$ for different values of ϕ and λ.

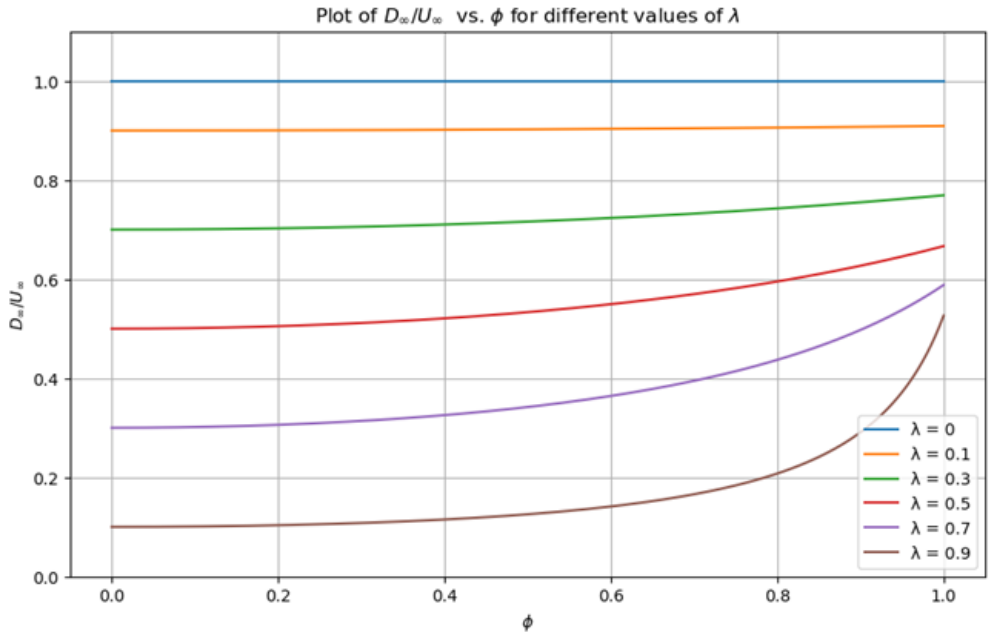


Figure 10: Comparative statics for disagreement/uncertainty ratio under sticky information

7 Discussion

At this point, it is useful to recap some of the key results that we have derived, especially in terms of the similarities and differences in the disagreement/uncertainty ratio under noisy and sticky information. Indeed, the steady-state disagreement/uncertainty ratio is upper-bounded by unity, and is decreasing in the degree of informational frictions (σ_ν^2 in the case of noisy information and λ in the case of sticky information). Moreover, the ratio is invariant with respect to shocks under noisy information but is time and state-dependent under sticky information (it is increasing in the absolute magnitude of shocks). Under both noisy and sticky information, the ratio is increasing in the degree of state persistence (ϕ). A salient difference, however, is that the steady-state disagreement/uncertainty ratio is increasing in fundamental volatility (σ_ϵ^2) under noisy information, but is invariant with respect to σ_ϵ^2 under sticky information.

To roughly fit the observed time series properties of the disagreement/uncertainty ratio (excluding the COVID-19 period, which appears to be somewhat anomalous), there are

two ways in which one could generate time variation in the disagreement/uncertainty ratio. The first approach would be to assume sticky information, whereby movement in the ratio would be driven by shocks. However, this explanation is somewhat unsatisfactory for three reasons. First, the disagreement/uncertainty ratio generally lies below one, and even on the few occasions where this does not hold, the ratio does not greatly exceed unity. Second, the periods in which the ratio exceeds unity appear to be times of high fundamental volatility, such as the Great Recession, an observation that is much more concordant with the noisy information than with the sticky information hypothesis. Under sticky information, it is large shocks relative to fundamental volatility that produce transitory variation in the disagreement/uncertainty ratio, which would imply that spikes in the ratio should occur at a homogeneous rate through time and that shocks to fundamental volatility should not affect the steady-state value of ratio. Third, the findings of Coibion and Gorodnichenko (2012) suggest that disagreement responds quite negligibly to shocks in general, again providing greater empirical support for the noisy information model relative to the sticky information model.

Instead, one could largely explain the time series properties of the disagreement/uncertainty ratio by assuming noisy information, coupled with stochastic volatility of fundamentals. It is rather difficult to analyze stochastic volatility in the fully dynamic setting, since the Kalman filter would no longer converge to a meaningful steady state, so we just discuss the intuition in the static case. Indeed, the empirical dynamics of core PCE inflation (we focus on this metric here as an empirical case study) are well described by a random walk, and assuming that forecasters are predicting the incremental changes in inflation relative to the previous year, the static model appears to be entirely reasonable in capturing this scenario. Since the disagreement/uncertainty ratio is increasing in σ_ϵ^2 and tends to 1 as fundamental volatility becomes large, the canonical noisy information model is a natural candidate to rationalize the observed spikes in the disagreement/uncertainty ratio during periods of heightened macroeconomic volatility.

One difficulty, however, lies in explaining the seasonal dynamics of the disagreement/uncertainty ratio. We would naturally expect that the noisiness of forecasters' signals diminishes over the four quarters of the year, which would imply that the disagreement/uncertainty ratio should *rise* between Q1 and Q4. In contrast, the precise opposite occurs in the data. Nonetheless, this limitation can be overcome by introducing a degree of *irreducible uncertainty*; that is, we posit that forecasters only obtain noisy signals about a latent state, where the true state is the latent state plus some independent noise. Denoting this irreducible variance by C, uncertainty now becomes $\frac{\sigma_\epsilon^2\sigma_\nu^2}{\sigma_\epsilon^2+\sigma_\nu^2} + C$, and the disagreement/uncertainty ratio need not be monotonically decreasing in σ_ν^2 anymore, as shown in Figure 11.

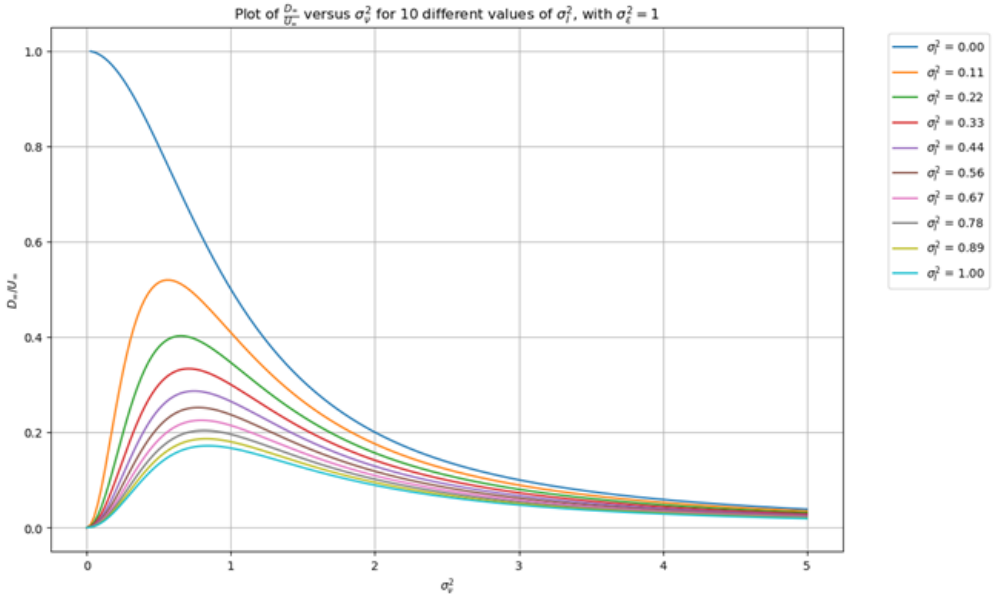


Figure 11: Non-monotonicity of disagreement/uncertainty ratio with irreducible uncertainty

While the idea of irreducible uncertainty has been relatively unexplored in previous literature relating to models of imperfect information in macroeconomics, it certainly appears to be a useful theoretical construct to capture residual uncertainty even when signals are perfectly revealing. However, this assumption does not come without loss, since the introduction of irreducible uncertainty implies that the disagreement/uncertainty ratio no longer tends to unity as fundamental volatility becomes large. To preserve the limiting behavior, we would require σ_ϵ^2 and σ_ν^2 to both increase, but the degree of increase in σ_ϵ^2 should be an

order of magnitude greater than that of σ_ν^2.

To illustrate these ideas, we perform a rough calibration of the static noisy information model to fit the observed time series of disagreement, uncertainty and the disagreement/uncertainty ratio for core PCE inflation, where the noise variance exhibits seasonal variation, with the identifying assumption that the noise variance is monotonically decreasing over the four quarters of the year. We essentially perform a simulated method of moments estimation protocol, where the objective function minimizes the sum of squared deviations between the simulated and empirical means of uncertainty and the disagreement/uncertainty ratio in each of the four quarters of the year. Indeed, the chosen parameters of $\sigma_{\nu,1}^2 = 0.24$, $\sigma_{\nu,2}^2 = 0.18$, $\sigma_{\nu,3}^2 = 0.08$, $\sigma_{\nu,4}^2 = 0.016$ (where the numerical subscripts denote the quarter under consideration) allow us to closely match the empirical moments. In Figures 12 and 13, we show the simulated time series of disagreement, uncertainty and the disagreement/uncertainty ratio, where the baseline parameters are overlaid with occasional random spikes in σ_ϵ^2 and σ_ν^2.

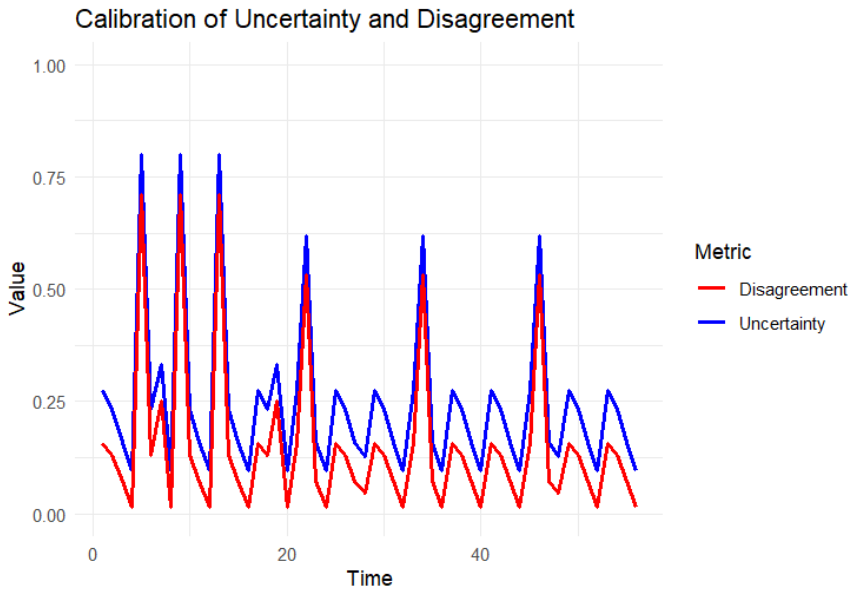


Figure 12: Calibrated dynamics of uncertainty and disagreement in noisy information model with stochastic volatility and irreducible uncertainty

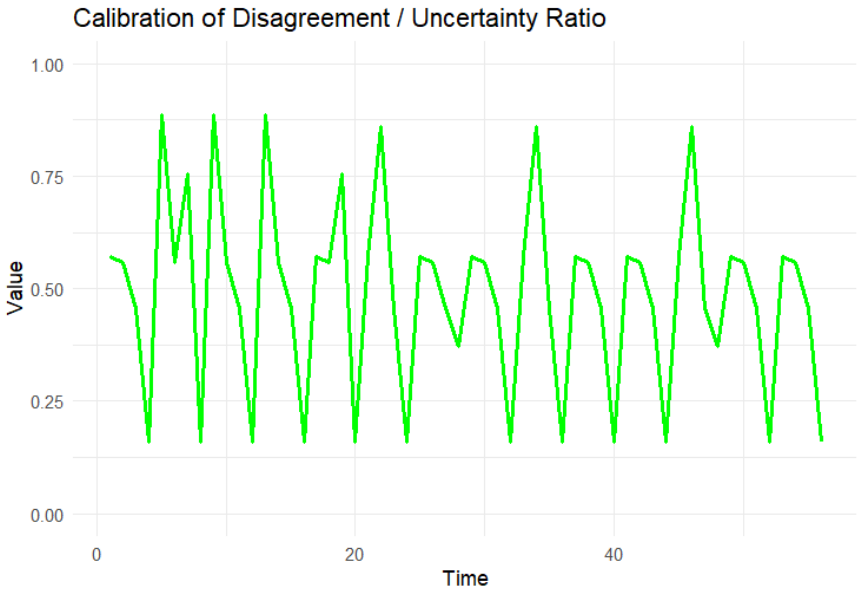


Figure 13: Calibrated dynamics of disagreement/uncertainty ratio in noisy information model with stochastic volatility and irreducible uncertainty

The primary issue, however, with the calibration is that very large increases in σ_ϵ^2 are required to yield increases in the disagreement/uncertainty ratio toward 1 (for illustrative purposes in generating the figures, we multiply σ_ϵ^2 by 50 and σ_ν^2 by 3 during designated "crisis periods"). This results in excessively large spikes in uncertainty relative to what is observed in the empirical data. Thus, the benefits of irreducible uncertainty in accounting for the seasonal dynamics comes at the cost of making it significantly more difficult to rationalize the spikes in the disagreement/uncertainty ratio.

As suggested by the earlier theoretical analysis, it is also worthwhile to examine the behavior of realized forecast errors in parallel with uncertainty and disagreement. It is interesting to note, as Figure 14 shows, that there are systematic deviations between implied uncertainty and realized forecast errors over the four quarters of the year (forecasters systematically underestimate their uncertainty in Q1 and overestimate their uncertainty in Q4). In essence, uncertainty implied by probability forecasts exhibits excess stickiness, perhaps due to asymmetric loss functions as in Capistrán and Timmermann (2009), and the observed residual uncertainty is at least in part an artifact of this phenomenon.

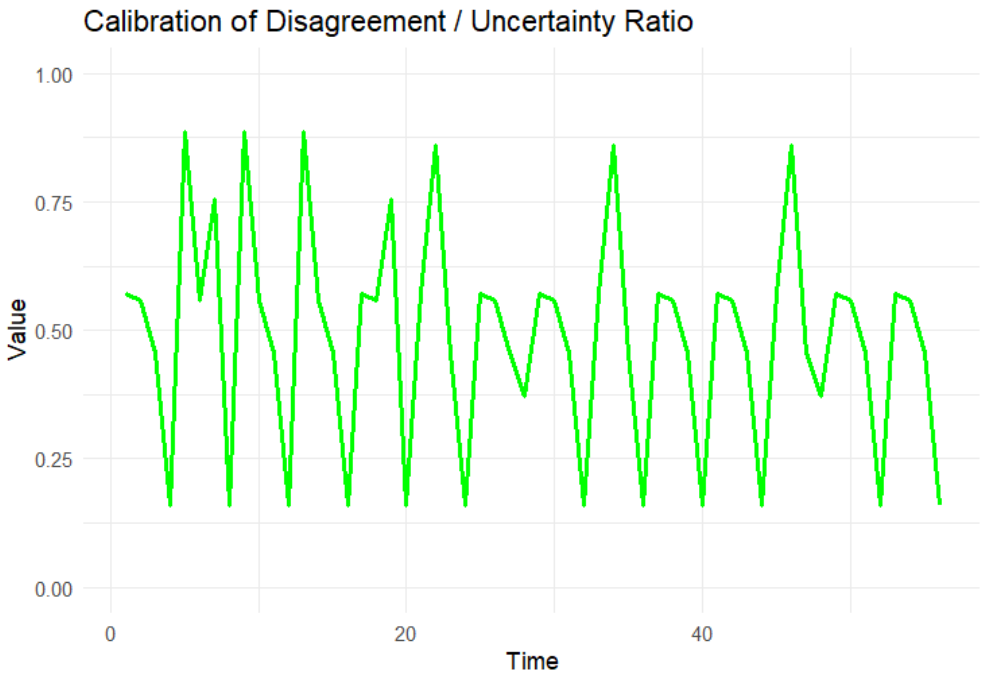


Figure 14: Disagreement, uncertainty and realized forecast errors for core PCE inflation

Naturally, it is likely that there are elements of truth to different models, and one example of a situation where it is difficult to rationalize either the noisy or sticky information models is during the post-COVID inflation. Here, the disagreement/uncertainty ratio is anomalously large, suggesting that other forces, such as differences in priors or perceived models across forecasters, may be playing an important role. Nonetheless, we concur with the finding in Coibion and Gorodnichenko (2012) and Coibion and Gorodnichenko (2015) that the noisy information model does well in capturing the behavior of forecasters during "normal times", though we also identify some caveats that necessitate modifications to the standard model.

8 Conclusions

While disagreement has been widely used as an empirical proxy for uncertainty, comparatively little is known about the theoretical soundness of this practice. In this paper, motivated by observed time series and cross-sectional properties of disagreement and uncertainty in forecasts taken from the Survey of Professional Forecasters, we derive a novel suite of

theoretical results to fill this void.

We first derive a general theorem that shows that under a common prior with rational expectations, expected disagreement cannot exceed expected uncertainty, even when information structures are arbitrary. We also show additional theoretical properties of disagreement, uncertainty and the disagreement/uncertainty ratio in setups based on the noisy information models of Woodford (2003) and Sims (2003), as well as the sticky information models in the style of Mankiw and Reis (2002). An important auxiliary contribution of our paper is to analytically characterize both disagreement and uncertainty under noisy information with state persistence by solving for fixed points of appropriate Riccati equations. We derive a rich set of comparative statics of the disagreement/uncertainty ratio under both informational regimes with respect to fundamental volatility, the extent of noise (in the case of noisy information), the degree of informational rigidity (in the case of sticky information) and the amount of persistence in the state variable.

Based on these theoretical results, we investigate the ability of the noisy and sticky information models to fit the observed dynamics of disagreement and uncertainty in data from the SPF. We concur with Coibion and Gorodnichenko (2012) and Coibion and Gorodnichenko (2015) that the noisy information model outperforms the sticky information model in fitting the data, but we show that some modifications to the standard model, namely, stochastic volatility and irreducible uncertainty, are required to reproduce some of the observed empirical patterns. Indeed, even with these modifications, some challenges remain in jointly fitting the seasonal patterns in uncertainty and disagreement and the rise in the disagreement/uncertainty ratio during periods of elevated volatility, a phenomenon that appears to be partly attributable to "excess stickiness" in subjective uncertainty.

A simultaneous benefit and limitation of the analysis in this paper is the focus on common prior environments. While this assumption permits analytical tractability and indeed appears to be a valuable empirical benchmark in most situations, it is certainly true that heterogeneity in priors and models cannot always be neglected. A natural extension of the

analysis in this paper would be to consider a broader range of informational frictions of this type, though it is quite likely that clear-cut theoretical results of the kind derived here will be significantly more difficult to obtain. We leave this interesting and important avenue for future work.

References

Barron, O. E., O. Kim, S. C. Lim, and D. E. Stevens (1998): "Using analysts' forecasts to measure properties of analysts' information environment," *Accounting Review*, 421–433.

Barron, O. E. and P. S. Stuerke (1998): "Dispersion in analysts' earnings forecasts as a measure of uncertainty," *Journal of Accounting, Auditing & Finance*, 13, 245–270.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020): "Overreaction in macroeconomic expectations," *American Economic Review*, 110, 2748–2782.

Capistrán, C. and A. Timmermann (2009): "Disagreement and biases in inflation expectations," *Journal of Money, Credit and Banking*, 41, 365–396.

Coibion, O. and Y. Gorodnichenko (2012): "What can survey forecasts tell us about information rigidities?" *Journal of Political Economy*, 120, 116–159.

——— (2015): "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 105, 2644–2678.

D'Amico, S. and A. Orphanides (2008): "Uncertainty and disagreement in economic forecasting," .

Diether, K. B., C. J. Malloy, and A. Scherbina (2002): "Differences of opinion and the cross section of stock returns," *The Journal of Finance*, 57, 2113–2141.

Farmer, L. E., E. Nakamura, and J. Steinsson (2024): "Learning about the long run," *Journal of Political Economy*, 132, 3334–3377.

Gaur, V., S. Kesavan, A. Raman, and M. L. Fisher (2007): "Estimating demand uncertainty using judgmental forecasts," *Manufacturing & Service Operations Management*, 9, 480–491.

Giordani, P. and P. Söderlind (2003): "Inflation forecast uncertainty," *European Economic Review*, 47, 1037–1059.

Hamilton, J. D. (1994): *Time Series Analysis*, Princeton University Press.

Lahiri, K. and F. Liu (2006): "Modelling multi-period inflation uncertainty using a panel of density forecasts," *Journal of Applied Econometrics*, 21, 1199–1219.

Lahiri, K. and X. Sheng (2010): "Measuring forecast uncertainty by disagreement: The missing link," *Journal of Applied Econometrics*, 25, 514–538.

Lucas Jr, R. E. (1972): "Expectations and the Neutrality of Money," *Journal of economic theory*, 4, 103–124.

Mankiw, N. G. and R. Reis (2002): "Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve," *The Quarterly Journal of Economics*, 117,

1295–1328.

Patton, A. J. and A. Timmermann (2010): "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion," *Journal of Monetary Economics*, 57, 803–820.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50, 665–690.

Woodford, M. (2003): "Imperfect Common Knowledge and the Effects of Monetary Policy," in *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honour of Edmund S. Phelps*, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, Princeton: Princeton University Press, 25–58.

Zarnowitz, V. and L. A. Lambros (1987): "Consensus and uncertainty in economic prediction," *Journal of Political Economy*, 95, 591–621.

Zhang, X. F. (2006): "Information uncertainty and stock returns," *The Journal of Finance*, 61, 105–137.