Heterogeneous Expectations among Professional Forecasters

Christian Conrad

Kajal Lahiri

Heterogeneous Expectations among

Professional Forecasters

Christian Conrad[1] and Kajal Lahiri[2]

March 25, 2024

Abstract[3]

Macroeconomic expectations of various economic agents are characterized by substantial cross-sectional heterogeneity. This chapter focuses on heterogeneity in the expectations among professional forecasters, first presenting stylized facts and discussing theoretical explanations for heterogeneous expectations, before providing an overview of the empirical evidence supporting the different theories and pointing to possible directions for future research. A literature review is complemented by empirical evidence based on the Leibniz Centre for European Economic Research (ZEW) Financial Market Survey, covering the behavior of expectations heterogeneity during the recent surge in inflation in 2021 and 2022. A central finding is that differences in perceptions about the workings of the economy and heterogeneity in perceptions of the precision of new signals drive disagreement among professional forecasters. While the level of disagreement varies over the business cycle, differences in beliefs persist over time.

Keywords: disagreement, expectations, forecasts, rationality, survey data
JEL Classification: C53, D83, D84, E17, E37

---

[1] christian.conrad@awi.uni-heidelberg.de, Heidelberg University, ORCID: 0000-0001-7981-2615
[2] klahiri@albany.edu, University at Albany, SUNY, ORCID: 0000-0003-4767-8591

Introduction

Forecasts or expectations of economic variables inform the decision-making of policymakers, firms, and households. For example, central banks base their policies on inflation projections, fiscal planning depends on predictions of future revenues, and financial decision-making requires forming expectations about risks and returns of alternative assets (see also Chapters 12, 13, and 16 in this volume).

In this survey, we focus on the expectations of professional forecasters for macroeconomic variables and the empirical evidence on their properties based on survey data.[4] Despite substantial efforts to improve prediction accuracy for better-informed decision-making, empirical evidence shows persistent cross-sectional heterogeneity in the expectations of professional forecasters (for early evidence, see Zarnowitz and Lambros 1987; Mankiw et al. 2003). We address the question: Why do professional forecasters disagree?

The literature on econometric forecasting has introduced the concept of forecast optimality or rationality (e.g., Theil 1966; Granger 1969; Diebold and Lopez 1996; Elliott et al. 2005; Patton and Timmermann 2007). Based on the properties of optimal or rational forecasts, methods for testing forecast rationality have been developed and applied to the forecasts of professional forecasters (see, for example, Mincer and Zarnowitz 1969; Zarnowitz 1985; Nordhaus 1987). Forecast biases or inefficient usage of available information can lead to a rejection of forecast rationality and, hence, explain why forecasters disagree. However, the optimal prediction depends on the forecaster's loss function, i.e., on how a forecaster assesses forecast errors. As emphasized by Elliott et al. (2005, 2008), tests of forecast rationality are conditional on the

---

[4] For a recent survey on the expectations of households and firms, see Weber et al. (2022) and Chapter 7 of this handbook.

assumed loss function, and forecasts that might appear to violate rationality when assuming a symmetric loss function can be consistent with rationality when assuming an asymmetric loss function. When forecasters are heterogeneous regarding loss functions, disagreement among forecasters can arise even if all forecasters are rational. Alternative explanations for forecast heterogeneity are differences in prior beliefs about the data-generating process (e.g., about long-run output growth), different interpretations of the same publicly available information (see, for example, Lahiri and Sheng 2008), or different information sets (Patton and Timmermann 2010). Again, heterogeneity in expectations may arise even if all forecasters rationally update their expectations.

In general, rational forecasters have homogeneous expectations if they (i) are homogenous regarding loss functions, (ii) know the data-generating process, i.e., the structure of the economy, and (iii) have access to the same information.

For a long time, the prevalence in macroeconomic modeling has been to assume that economic agents, particularly professional forecasters, have *rational expectations* in the sense of Muth (1961). Under rational expectations, in knowing the economy's structure, agents form *forward-looking* expectations, making efficient use of available information. Rational expectations can be seen as a response to the critique of the shortcomings of earlier, backward-looking models of expectation formation. It paved the way for the so-called *rational expectations revolution* (see Lucas 1972, 1976, and Chapter 2 in this volume).

Despite the importance of rational expectations in economic modeling, macroeconomists have often been skeptical concerning the credibility of survey-based expectations data (Manski 2017). This has changed in recent years, and evidence from survey data for deviations from rational

expectations and for heterogeneity in expectations, which may have important implications for the transmission of monetary and fiscal shocks, has spurred interest in theoretical models that can generate such heterogeneity. Empirical approaches for testing deviations from rational expectations often rely on econometric tests of forecast optimality.

In the macroeconomic literature, heterogeneity in expectations has been mainly explained by focusing on deviations from the *full information rational expectations* (FIRE) framework (see Coibion and Gorodnichenko 2012 and Chapter 10 in this volume). First, heterogeneous expectations can be due to differences in information sets (Lucas 1973). Models with sticky (Mankiw and Reis 2002) or noisy (Sims 2003) information can generate forecast heterogeneity while preserving the rational expectations assumption. The empirical observation that individual forecasters tend to overreact to new information is inconsistent with such models and has led to models that allow for deviations from rationality. For example, in the diagnostic beliefs model of Bordalo et al. (2020), forecasters place too much weight on the most recent information. In other models, heterogeneity is driven by differences in subjective historical experiences (Malmendier and Nagel 2011, 2016) or by agents' subjective models of the economy (Andre et al. 2022).

As our summary of the empirical evidence shows, one of the primary sources of disagreement among professional forecasters is heterogeneity in prior beliefs, particularly concerning long-run outcomes. Disagreement in prior beliefs about long-run outcomes has been documented for various sample periods and can be attributed, at least partially, to forecasters' diverse historical and anthropological backgrounds.

This chapter is organized as follows. In an initial section, we present stylized facts documented in the literature concerning the expectations of professional forecasters. Subsequently, we discuss econometric and macroeconomic models that can generate and explain heterogeneous expectations, before presenting and relating the empirical evidence regarding the different models. Finally, we summarize the main findings and point to directions for future research.

Heterogeneity: Stylized Facts

Empirical evidence in the previous literature has been based on surveys such as the US Survey of Professional Forecasters (SPF), the Blue Chip survey, or Consensus Economics. These surveys ask professional forecasters about their expectations about various economic variables such as output growth and inflation. The SPF has been conducted quarterly since 1968, and the Blue Chip survey and Consensus Economics monthly since 1976 and 1989 respectively. While the SPF asks for predictions of US macroeconomic variables only, the Blue Chip survey and Consensus Economics ask for predictions of variables in various countries. Predictions are elicited as point forecasts and/or probabilistic forecasts. Participants are asked for either fixed-horizon (e.g., inflation two-quarters-ahead) or fixed-event (e.g., inflation in the current year) forecasts in each survey round. Fixed-event forecasts are particularly useful for the analysis of interpersonal forecast heterogeneity. For encompassing surveys on survey expectations and their econometric analysis, see Pesaran and Weale (2006) and Clements (2019).

We complement the literature review with empirical evidence from the ZEW Financial Market Survey (henceforth ZEW survey), which is conducted by the ZEW in Mannheim, Germany. Since November 2014, this survey has asked professional forecasters from banks and insurance companies for their expectations of Harmonised Index of Consumer Prices (HICP) inflation in the Eurozone. In each quarter (specifically, in February, May, August, and November),

participants are asked about their predictions of this year's inflation, next year's inflation, and inflation in two years' time. Thus, predictions of Eurozone inflation in a specific year are collected for each forecaster in the survey for forecast horizons from twelve-quarters-ahead to one-quarter-ahead. For details on the ZEW survey, see Brückbauer and Schröder (2022).
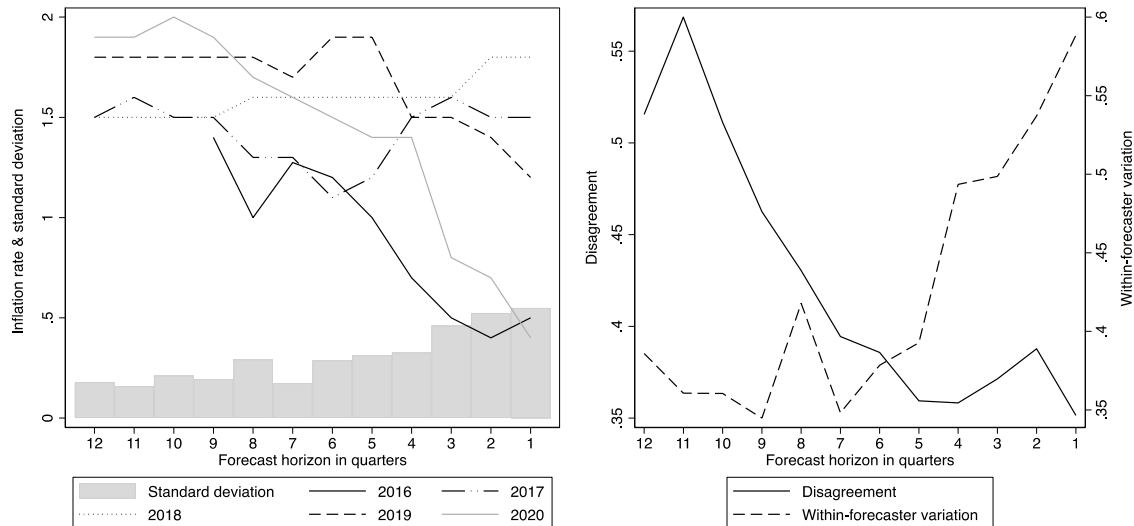

Figure 1: Left panel: Consensus forecasts of HICP inflation in the Eurozone from the ZEW Financial Market Survey. The lines show the fixed-event inflation forecasts for the target years 2016 to 2020. The black solid line begins for h=9 because the ZEW survey asked for a forecast of inflation in target year 2016 for the first time in November 2014. Grey bars indicate the standard deviation of the consensus forecasts at each forecast horizon. Right Panel: Disagreement and within-forecaster variation of individual forecasts of HIPC inflation in the Eurozone. The figure shows (the square root of) disagreement (left axis) and (the square root of) within-forecaster variation (right axis) at each forecast horizon. On average, the consensus forecast, disagreement, and within-forecaster variation are computed based on individual predictions of 146 survey participants.

The left panel in Figure 1 shows the *consensus forecasts* of the ZEW survey's participants for the years 2016 to 2020. Consensus forecasts are defined as the cross-sectional average of the individual point predictions concerning a specific forecast horizon when the survey is conducted (Zarnowitz and Lambros 1987). Each line shows how the consensus forecast of Eurozone inflation in a specific target year changed with the forecast horizon (e.g., the gray line illustrates how the consensus forecast of inflation in 2020 varied over time). The most extended forecast horizon is twelve quarters, and the shortest is one quarter. The figure shows that, independently of the target year, the consensus forecasts are close to the ECB's inflation target of (close to but

below) two percent when the forecast horizon is long. When the forecast horizon decreases and forecasters update their predictions, the consensus forecasts become more dispersed. This is illustrated by the gray bars, which show the standard deviation of the consensus forecasts at each forecast horizon. Intuitively, this is the behavior that we would expect from the forecasts when inflation follows a stationary process with an unconditional mean that equals the inflation target of the ECB.

Disagreement

While the left panel of Figure 1 illustrates how the consensus forecasts vary over time, it is not informative about forecast heterogeneity across survey participants. At a specific point in time, the heterogeneity or disagreement among forecasters concerning the *h*-step-ahead (e.g., two-quarters-ahead) forecast is commonly measured by the cross-sectional variance of the individual point predictions (see, e.g., Zarnowitz and Lambros 1987; Lahiri and Sheng 2008). The *h*-step-ahead disagreement is obtained by averaging the corresponding disagreements over time. In the following, we illustrate how disagreement varies with the forecast horizon and over the business cycle.

Forecast Horizon

Strong empirical evidence exists that disagreement among forecasters varies with the forecast horizon (Lahiri and Sheng 2008; Patton and Timmermann 2010). For fixed-event forecasts, disagreement is typically highest at the most extended forecast horizons and tends to decrease when the forecast horizon gets shorter. Nevertheless, even at short forecast horizons, considerable disagreement remains. The solid line in the right panel of Figure 1 illustrates this stylized fact for the inflation forecasts from the ZEW survey (again based on data for the target years 2016 to 2020).

An alternative perspective on the behavior of the individual forecasts can be gained by decomposing the total sum of squares of forecasts at each forecast horizon into the within-forecaster variation and the between-forecaster variation. The former is measured by first computing the time-variation of each forecaster's h-step-ahead forecasts and then averaging over forecasters. The latter is calculated as the cross-sectional variance of the forecaster-specific time-averaged h-step-ahead forecasts. The dashed line in the right panel of Figure 1 shows that within-forecaster variation is usually the lowest at the most extended forecast horizons and increases with decreasing forecast horizons. Interestingly, within-forecaster variation increases at forecast horizons of eight and four quarters respectively. This is when information about inflation two years and one year ahead of the target year becomes available.

As we will discuss below, strong disagreement in combination with low within-forecaster variation in long-term predictions can be explained by persistent differences in subjective beliefs about how the economy works. Disagreement in short-term predictions is likely due to different interpretations of the same information.

Persistence

Evidence shows forecasters persistently deviate from the consensus forecast on the upper or lower side (Batchelor 2007; Patton and Timmermann 2010; Boero et al. 2015). That is, some forecasters are persistently optimistic, while others are pessimistic. For each forecast horizon, the left panel of Figure 2 shows the time-averaged position of four selected participants of the ZEW survey in the cross-sectional forecast distribution. All four participants are characterized by persistent optimism or pessimism. As we discuss in "Importance of Prior Beliefs" below,

8

expectation persistence may be due to the "anchoring effect" of prior beliefs (Zellner 2002; Lahiri and Sheng 2010a).
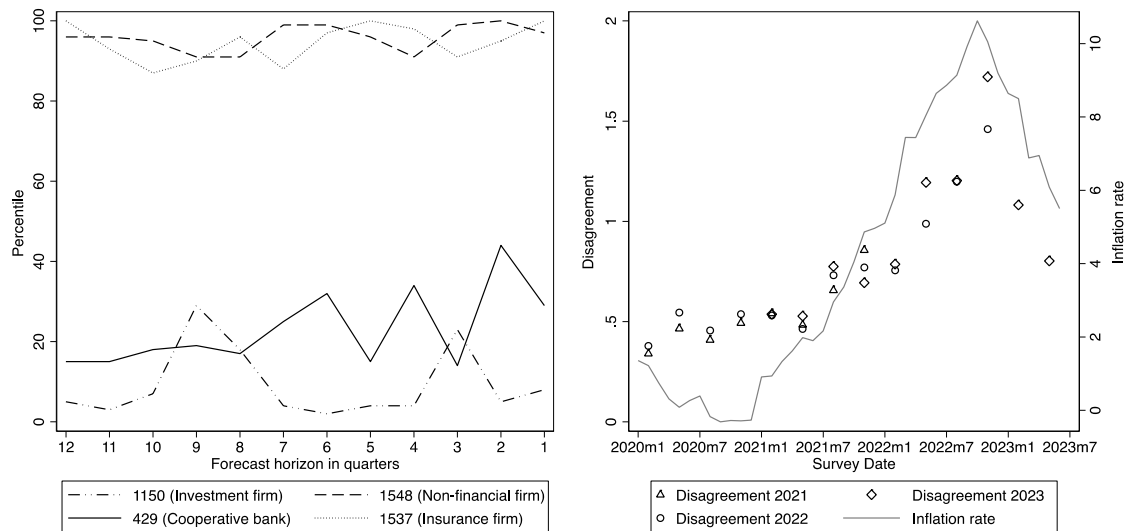

Figure 2: Left panel: For each forecast horizon, the figure shows the relative position (time average) of four selected forecasters in the cross-sectional forecast distribution. Right panel: Evolution of disagreement during the inflation surge in 2021 and 2022. The triangles, circles, and diamonds show (the square root of) disagreement among forecasters concerning inflation in 2021, 2022, and 2023. The solid line depicts actual HICP inflation in the Eurozone.

Time-Variation

There is also considerable variation in disagreement over time. For example, Mankiw et al. (2003) show that disagreement about inflation increases with the level of inflation and, in particular, when inflation changes substantially. Patton and Timmermann (2010) confirm the positive correlation between the level of inflation and disagreement. In addition, they provide evidence for a negative correlation between output growth and disagreement about future output growth, i.e., disagreement about output growth behaves counter-cyclically. Both findings are confirmed by Dovern et al. (2012). The right panel of Figure 2 shows how disagreement among the participants of the ZEW survey about inflation in 2021 (triangles), 2022 (circles), and 2023 (diamonds) evolved over time. Forecasts for those years were not included in Figure 1 because of the extreme surge in actual inflation. The right panel of Figure 2 illustrates that disagreement about inflation in 2021 and 2022 increased almost parallel with the level of the actual inflation

9

rate. In particular, disagreement jumped to higher levels than usually observed for even the most extended forecast horizons. During the 2021–2022 period, characterized by high uncertainty about future inflation, disagreement about inflation in 2021 and 2022 increased with decreasing forecast horizons and peaked in November 2021 and November 2022 respectively. Disagreement about inflation in 2023 started to fall with declining inflation rates in 2023.

Uncertainty vs. Disagreement

As many surveys (e.g., the SPF) do not only ask for point predictions but also probabilistic predictions, i.e., histogram forecasts, ex ante forecast uncertainty can be estimated at an individual level. This is typically done by assuming a parametric distribution and fitting this distribution to the respondent's histogram forecast (Giordani and Söderlind 2003; Engelberg et al. 2009). The fitted distribution's estimated standard deviation is then used to measure individual ex ante uncertainty. Similar to the findings reported in the previous section, there is strong evidence for heterogeneity in individual uncertainty and evidence for persistence in the relative level of individual uncertainties, i.e., some forecasters are persistently more/less uncertain than others (Boero et al. 2015; Rich and Tracy 2021). Thus, there is not only disagreement in point predictions but also in individual forecast uncertainties.

Similar to the construction of the consensus forecast, the "typical" uncertainty can be measured by the average individual uncertainty (Lahiri and Sheng 2010b). Because not all surveys elicit probabilistic predictions, it has been common to use disagreement among forecasters as a proxy for the average individual uncertainty. However, as pointed out by Zarnowitz and Lambros (1987), disagreement might be low or even zero, i.e., all forecasters make the same point prediction, while average individual uncertainty can be high. Conversely, each forecaster might be very confident about their prediction, i.e., average individual uncertainty is low, but

disagreement can be high. Lahiri and Sheng (2010b) provide a framework for linking disagreement to average individual uncertainty. In their model, forecasters receive and optimally combine a private and a public signal. Individual uncertainty can be decomposed into the perceived uncertainty of aggregate shocks and the variance of idiosyncratic shocks. In this setting, average individual uncertainty can be written as disagreement plus the perceived uncertainty of aggregate shocks. Thus, disagreement is likely to be a good proxy for uncertainty only during periods in which the perceived uncertainty of aggregate shocks is low. Indeed, the empirical evidence suggests that the link between disagreement and uncertainty is relatively weak and depends on the sample period used (Glas 2020).

Two Perspectives on Forecast Heterogeneity

This section presents theoretical explanations for forecast heterogeneity. The first perspective is motivated by the econometric literature on forecast evaluation and tests for forecast rationality. The second perspective comes from macroeconomic modeling. Research on forecast heterogeneity is currently striving, and the econometric tests and macroeconomic approaches we present are only a selection of available models motivated by our perception of the literature.

Forecast Rationality—The Econometric Perspective

The econometric literature on forecast rationality typically takes the loss function of a forecaster as a starting point. Based on the loss function, the data-generating process, and available information, the optimal forecast can be determined, and properties of the optimal forecast and the corresponding forecast errors and revisions can be derived. In the following, we will assume that the process to be forecasted is covariance-stationary. In particular, this implies that the process has a time-invariant mean and variance. First, "Squared Error Loss" will focus on the situation where a forecaster has a squared error (SE) loss function. In this case, a rational forecast

is unbiased; hence, forecast errors have zero mean. However, this does not need to be the case. The optimal forecast will typically be biased if forecasters have an asymmetric loss function, as discussed subsequently ("Asymmetric Loss Functions"). The role of prior beliefs in forming expectations is then highlighted in "Importance of Prior Beliefs."

Squared Error Loss

It is common to assume that forecasters minimize the SE loss conditional on a specific information set. The information set contains past observations of the process to be forecasted and, potentially, other variables that are observed when the forecast is made. In this setting, the optimal forecast is given by the conditional mean (see, e.g., Diebold and Lopez 1996; Patton and Timmermann 2007). Tests of forecast rationality are based on or derived from the properties of the forecast errors.

First, the optimal forecast is unbiased, i.e., the forecast errors have a mean of zero. A simple test of forecast rationality can be conducted by running a regression of the forecast error on a constant and testing whether the constant is significantly different from zero. Because the process to be forecasted is covariance-stationary, the h-step-ahead forecast should converge to the unconditional mean of the process. Thus, disagreement among forecasters about long-run predictions suggests differences in subjective opinions about the unconditional mean of the process.

Second, the forecast error should be uncorrelated with all variables in the information set. Hence, tests for forecast rationality check whether a forecaster efficiently uses all available information. Since it is often unclear which variables were included in an information set that a forecaster was conditioning on, obvious choices for those variables are forecasts or lagged forecast errors. For

example, testing whether the forecast has explanatory power for the future forecast error is common. A version of this test suggested by Mincer and Zarnowitz (1969) is a regression of the outcome on a constant and the forecast. This test of forecast rationality has the joint null hypothesis that the constant and the slope are equal to zero and one, respectively. An alternative test of forecast rationality was suggested by Nordhaus (1987). His test is designed for settings where fixed-event forecasts are available. It is based on a regression of the *current h*-step-ahead forecast error on *past* forecast revisions for the same target variable. If forecasters use available information efficiently, past forecast revisions should not predict the current forecast error. We will refer to this approach as the "Nordhaus test" in the following.

Third, the variance of the forecast error or, alternatively, the forecast's mean squared error (MSE) should be a non-decreasing function of the forecast horizon (e.g., Patton and Timmermann 2007). Again, because of covariance-stationarity, the variance of the forecast error will converge to the unconditional variance of the process as the forecast horizon increases. Isiklar and Lahiri (2007) focus on the difference of the MSE at forecast horizons *h*+1 and *h*. This difference should be non-negative for the optimal forecast and can be considered a measure of the new information content that becomes available at forecast horizon *h*. If forecasters deviate from the optimal forecast, i.e., do not efficiently use available information, the change in the MSE is the sum of two components. The first component corresponds to the MSE change observed when information is efficiently used. The second term reflects the price the forecaster has to pay for deviating from efficiency. Isiklar and Lahiri (2007) suggest comparing the change in the MSE to the expectation of the squared forecast revision from *h*+1 to *h*. For the optimal forecast, the expected squared revision equals the change in the MSE. If the expected squared revision is *smaller/larger* than the change in the MSE, the forecaster tends to *underreact/overreact* to new information.

Asymmetric Loss Functions

Forecasters do not necessarily have a SE loss function. Instead, it is often reasonable to assume that forecasters have an asymmetric loss function, e.g., due to the underlying decision problem and strategic or psychological causes (Capistrán and Timmermann 2009). For example, it might be costlier for a central bank to underestimate inflation than to overestimate it. If the costs of under- and overpredicting are no longer symmetrical, the optimal prediction typically deviates from the conditional mean (e.g., Granger 1969; Patton and Timmermann 2007). For example, under a Linex loss function, the optimal forecast is the conditional mean plus a term that depends on an asymmetry parameter and the conditional variance of the target variable (Zellner 1986). This implies that the optimal forecast is no longer unbiased and that even the long-run forecast deviates from the unconditional mean. In addition, even one-step-ahead forecast errors can be serially correlated if the conditional variance is persistent. Thus, a rejection of forecast rationality based on the tests described in "Squared Error Loss" above may be spurious due to an asymmetric loss function (Elliott et al. 2005, 2008). However, standard tests may be adjusted, for example, by regressing the forecast error on the conditional variance and other variables that are observable when the forecast is made. Controlling for the conditional variance, those other variables should still have no predictive power for the forecast errors (Pesaran and Weale 2006). Asymmetric loss functions can explain disagreement in point predictions; even if all forecasters base their predictions on the same information set, there will be heterogeneity in predictions due to different asymmetry parameters. Under a Linex loss function, disagreement among forecasters will vary with the level of uncertainty concerning the target variable (Capistrán and Timmermann 2009). This loss function can also explain the persistence in the relative ranking of the individual forecasts described earlier. If all forecasters use the same conditional mean and conditional

variance forecasts, the relative ranking of the forecasts is determined by the individual asymmetry parameters.

Importance of Prior Beliefs

At the beginning of this chapter, we documented that forecast disagreement varies with the forecast horizon. Lahiri and Sheng (2008, 2010a) use a Bayesian learning model to explain this behavior. In their model, forecasters have prior beliefs about the outcome variable, and, in addition, can interpret public signals differently; in each period, beliefs are updated according to Bayes' rule. In this model, disagreement among forecasters is generated by three components: prior beliefs, the weight attached to prior beliefs (depending on the precision of both the initial belief and the public signal), and the interpretation of public signals. Specifically, the Bayesian learning model provides a means of studying the relative importance of prior beliefs and heterogeneity in incorporating new, publicly available information for explaining heterogeneity in forecasters' predictions at different forecast horizons. Notably, prior beliefs induce stickiness in expectations so that stickiness is consistent with rational updating of forecasts. Lahiri and Sheng (2008) propose a forecast horizon-specific regression of the *current* forecast error on the *current* forecast revision to test the rationality of forecasters at an individual level. It is important to note that the Lahiri and Sheng (2008) regression differs from the "Nordhaus test" because the regressor is the *current* and not the *past* forecast revision. Thus, the Lahiri and Sheng (2008) regression tests how *new* information is incorporated. Similar to Isiklar and Lahiri (2007), a *positive/negative* slope coefficient implies *under-/overreaction* to public news.

Patton and Timmermann (2010) also distinguish between information signals and prior beliefs or subjective models. They assume that the variable to be forecasted is the sum of a persistent and a transitory component. The forecasters' observation is contaminated with common and

idiosyncratic shocks. First, assuming a SE loss function, forecasters compute the optimal prediction, conditional on available information. Although all forecasters use the same model, there is forecast heterogeneity due to differences in information sets. Second, Patton and Timmermann (2010) introduce an additional layer of heterogeneity: forecasters have subjective beliefs about the unconditional mean of the variable to be forecasted and shrink the optimal forecast towards the subjective belief. Thus, similarly to Lahiri and Sheng (2008), disagreement is strongly driven by heterogeneity in prior/subjective beliefs. Giacomini et al. (2020) combine heterogeneous priors and differences in how new public information is interpreted (due to heterogeneous statistical models) with inattentiveness, i.e., infrequent updating, to explain why forecasters "agree to disagree."

Expectation Formation in Macroeconomic Models

The macroeconomic literature has recently collected ample empirical evidence against the FIRE paradigm. In this section, we first discuss models that deviate from the assumption of full information and explain how those deviations generate disagreement. We then turn to models that allow for non-rational expectations and introduce complementary approaches that focus on experiences and subjective models of the economy.

Sticky and Noisy Information Models

Sticky information models (see Mankiw and Reis 2002) and noisy information models (see Sims 2003) have been suggested to introduce information rigidities. In both models, forecasters are rational (given available information) and minimize the SE loss. In sticky information models, updating information sets is costly, and disagreement is generated because, in each period, only a fraction of the forecasters updates information. In the noisy information model, forecasters observe noisy public and private information, and the amount of disagreement depends on the

variance of the innovation to the private signal. Coibion and Gorodnichenko (2012) show that noisy and sticky information models can be distinguished by the predicted effects of economic shocks on disagreement. While economic shocks should not affect disagreement in the noisy information model, disagreement responds to shocks in the sticky information model. While agents act rationally at the individual level, both models imply that the consensus forecast underreacts to new information. Coibion and Gorodnichenko (2015) suggest regressing *current* forecast errors on *current* forecast revisions to test this property. Under the null hypothesis of FIRE, the forecast errors of the consensus forecast are unpredictable. Under the alternative, both sticky and noisy information models imply that the coefficient on the forecast revision is positive. However, it is essential to note that under both models (sticky and noisy), individual forecast errors should not be predictable based on individual forecast revisions. The regression proposed by Coibion and Gorodnichenko (2015) is the same as the Lahiri and Sheng (2008) regression but uses the consensus forecast instead of the individual forecasts. Their important contribution is the insight that this regression can be used to test sticky and noisy information models against the FIRE paradigm.[5]

Deviations from Rationality

Recently, several models of expectation formation that allow for deviations from rational expectations have been introduced. A prominent example is the diagnostic expectations model proposed by Bordalo et al. (2020). In this model, agents make a forecast that consists of the rational forecast plus a component that overweighs the most recent information. Forecast updates in response to good/bad news are too optimistic/pessimistic. A testable prediction of this model is that individual forecasters tend to overreact to new information. Interestingly, this overreaction

---

[5] For other approaches that confront sticky and noisy information models with survey data and link them to disagreement, see, for example, Andrade and Le Bihan (2013), Dovern (2015), Andrade et al. (2016), and Dovern and Hartmann (2017).

at the individual level is consistent with the underreaction of the consensus forecast. This prediction distinguishes the model from noisy and sticky information models, which, as discussed before, imply forecast rationality at the individual level.

Subsequently, several approaches for more granular modeling of under-/overreaction have been proposed. For example, Angeletos et al. (2021) suggest a model that combines noisy information and overextrapolation. Their model can explain why forecasters initially underreact to news but overreact subsequently. Based on evidence from a randomized lab experiment, Afrouzi et al. (2023) develop a model where overreaction to news depends on the forecasted process properties and the forecast horizon. Overreaction is more pronounced for more transitory processes and at longer forecast horizons. In their model, recent observations affect the forecasters' beliefs about the long-run mean of the process to be forecasted.

Experiences, Subjective Models, and Narratives

Malmendier and Nagel (2011, 2016) show that individual lifetime experiences strongly affect how agents form expectations. For example, if agents have experienced bad stock market performance, this has a long-lasting negative effect on their expectations of future returns. Agents with different experiences have different expectations. This is true even for professional forecasters. For example, Malmendier et al. (2021) use a model of experience-based learning to explain heterogeneity in the inflation projections of the members of the Federal Open Market Committee by their individual inflation experiences. Similarly, Benchimol et al. (2022) show that experiences such as having worked at a central bank affect expectations. Conrad et al. (2022) find that experiences also interact with how agents process new information. For example, households that have experienced higher inflation rates are less responsive to news about inflation in the media. This has important policy implications. For example, as Conrad et al.

(2022) discussed, the expectations of households with higher inflation experience might be harder to "manage" by central banks (see also Chapter 12 in this volume).

Another source of disagreement are subjective models of the economy. If forecasters believe in different models of the economy, there will be heterogeneity in how they update their forecasts in response to the same new information. For example, Conrad et al. (2022) find that households that experienced high inflation in the past are more likely to update their inflation expectations upwards in response to a contractionary monetary policy shock. They argue that inflation experiences shape the economic models that households rely on when processing new information. Andre et al. (2022) show that even among professional forecasters, there is heterogeneity, for example, in how they update their unemployment forecast in response to a monetary policy shock. Their findings highlight the role of "associative memory" in explaining heterogeneity; forecasters might entertain several subjective models, and the context or individual experiences can determine which model the forecaster is using. Conrad et al. (2024) provide evidence for the effects of narratives on how professional forecasters update stock market forecasts in response to new information. Using the ZEW survey, they show that at the end of 2022, forecasters entertained different narratives about the development of inflation in Germany in 2023. For example, one group of forecasters believed in the narrative that inflation would stay high and further monetary tightening in combination with a recession would negatively impact the stock market. When provided with an information treatment concerning the future development of inflation, the narratives that forecasters entertained affected how they updated their stock market expectations.

Empirical Evidence

This section summarizes empirical evidence for and against the explanations for heterogeneity discussed in the previous section. We mainly focus on tests of forecast rationality.

Result 1: Biases at the Individual Level

The simplest explanation for disagreement is heterogeneity in individual biases. Because there is ample evidence for biases at the level of individual forecasters, we only refer to a few selected studies. Using data from the SPF for the 1968–1979 period, Zarnowitz (1985, 299) shows that "almost all forecasters underestimated inflation and did so increasingly for the more distant future." About half the participants have significant biases. Similarly, Davies and Lahiri (1995) provide evidence for biases in the Blue Chip survey's inflation and output growth expectations of individual forecasters. Using data from Consensus Economics, Batchelor (2007) finds evidence for systematic biases in individual expectations, particularly in output growth forecasts. Capistrán and Timmermann (2009) analyze the inflation predictions of the individual forecasters from the SPF. They provide evidence that at least half of the forecasters have a significant bias, whereby most forecasters tend to under-predict inflation. For further evidence of biases at the individual level, see, for example, Dovern and Weisser (2011) and Sheng (2015). Overall, the evidence suggests that biases depend on the forecast horizon. While individual forecasts are most biased at the longest forecast horizons (Zarnowitz 1985; Juodis and Kučinskas 2023), forecasts are essentially unbiased at very short horizons, i.e., shortly before a data release (e.g., Conrad et al. 2023).

Finally, it is essential to note that the evidence for biases is much stronger at the individual level than when using the consensus forecast, where individual biases tend to cancel out (Zarnowitz and Lambros 1987).

Result 2: Asymmetric Loss Function

As discussed earlier in this chapter, forecast biases may arise due to asymmetric loss functions. Elliott et al. (2008) find that the evidence against forecast rationality in SPF data based on the Mincer and Zarnowitz (1969) regression is driven by the joint hypothesis of rationality and a symmetric loss function. Instead, when allowing for asymmetric loss, the inflation and output growth forecasts are consistent with rationality. Using the latest weighting matrices for generalized method of moments estimation, Krüger and LaCrone (2019) show that the Elliott et al. (2008) approach leads to precise estimates of the degree of asymmetry of the loss function. Capistrán and Timmermann (2009) report that roughly half of the SPF survey respondents are characterized by a Linex loss function. Even when allowing for asymmetric loss, they find evidence for a time-invariant bias for more than thirty percent of the forecasters. This time-invariant bias, in combination with the asymmetric loss, helps to explain why forecasters tended to underestimate inflation before the Great Moderation and to overestimate inflation during that period. Before the Great Moderation, inflation volatility was high and, in combination with a negative asymmetry parameter, forecasts underestimated inflation despite a positive time-invariant bias. During the Great Moderation, inflation volatility was low, and the positive time-invariant bias dominated. Capistrán and Timmermann (2009) also find that the ranking of the forecasters' predictions is persistent over time, as predicted by the model, and that disagreement is positively related to the conditional variance of inflation. If the volatility of inflation is positively related to the level of inflation, as empirical evidence suggests (Conrad and Hartmann 2019), there is also a positive relation between disagreement and the level of inflation (as

observed in "Heterogeneity: Stylized Facts" above). Further evidence for asymmetric loss functions is provided in, for example, Capistrán (2008) and Patton and Timmermann (2007). As shown in Conrad and Hartmann (2023), disagreement or, if available, the average individual variance can be used to de-bias the consensus forecast when individual forecasters entertain asymmetric loss functions.

Result 3: Prior Beliefs Induce Persistence

As discussed in "Importance of Prior Beliefs" above, Lahiri and Sheng (2008, 2010a) investigate the determinates of heterogeneity in expectations at different forecast horizons. Using data from Census Economics, they show that prior beliefs and heterogeneity in interpreting public signals are the most critical drivers of heterogeneity in expectations. While prior beliefs dominate at long-forecast horizons, public signals are more important at shorter horizons. Patton and Timmermann (2010) obtained similar results but argued that prior beliefs are generally more important for explaining expectation heterogeneity than information signals. Their empirical results also suggest that the weight attached to prior beliefs increases in crisis periods, which explains heightened disagreement during those periods. Finally, differences in prior beliefs help to explain why differences in forecasts (optimism/pessimism) persist over time.

Result 4: Underreaction of the Consensus Forecast and Overreaction of Individual Forecasts

Empirical evidence suggests that forecasters do not efficiently use all available information. We first discuss evidence at the level of the consensus forecast and then for the individual forecasters.

The approach suggested by Isiklar and Lahiri (2007) and Lahiri (2012) can be used to check for forecast horizon-specific inefficiencies. Based on the Blue Chip consensus forecasts for output growth for 1986–2009, Lahiri (2012) shows that the change in the MSE is bigger than the average squared forecast revision, i.e., the consensus forecast underreacts to new information. The

underreaction is most pronounced at the middle forecast horizons of 11 to 15 months. Typically, at these horizons, forecasters revise their predictions most strongly (see Lahiri and Sheng 2010a). Similar evidence is provided by Isiklar and Lahiri (2007).

Using SPF data, Coibion and Gorodnichenko (2015) provide further evidence for the underreaction of the consensus forecast by regressing the forecast error on the current forecast revision. They interpret this as evidence against the FIRE hypothesis and as evidence in favor of models that drop the *full information* assumption. Combined with earlier empirical evidence in Coibion and Gorodnichenko (2012) that disagreement does not appear to respond to economic shocks, the data aligns with the noisy but not the sticky information model. In contrast, using forecast revisions of fixed target forecasts as a measure of shocks, Hur and Kim (2016) find that the dynamics of dispersion in survey forecasts are consistent with sticky information models as opposed to conventional noisy information models.

In Figure 3, we reproduce Lahiri's (2012) analysis for the consensus inflation forecasts from the ZEW survey. First, the figure shows how forecast accuracy (as measured by the mean squared error, MSE) increases with decreasing forecast horizons. For example, substantial gains in forecast accuracy materialize at forecast horizons of eight and four quarters respectively, i.e., when realized inflation two and one years before the forecast target year becomes available. By comparing the change in the MSE (i.e., the MSE difference) and the mean squared forecast revision (MSR), it becomes evident that the consensus forecast underreacts to new information at most forecast horizons. The underreaction is most pronounced at forecast horizons of three, four, and eight quarters.
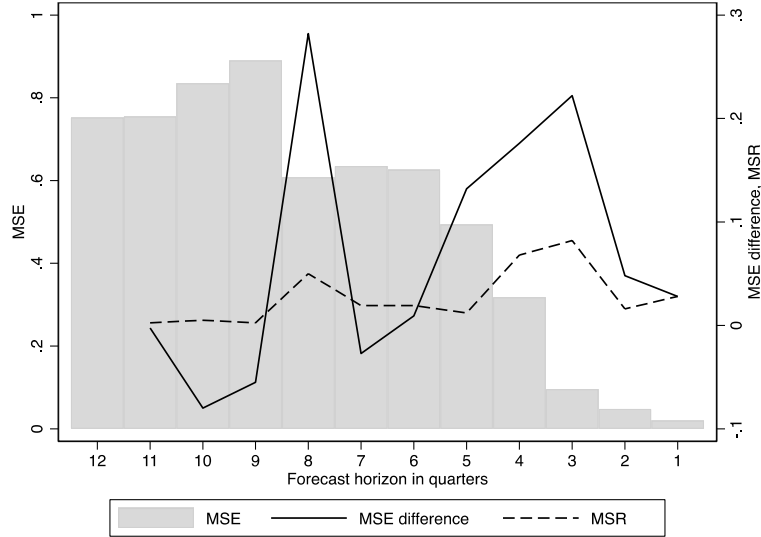
Figure 3: The grey bars show the consensus forecast's horizon-specific mean squared error (MSE). At each forecast horizon h, the solid line shows the difference between the (h+1)-step MSE and the h-step MSE (denoted by MSE difference). The dashed line shows the mean squared revision (MSR), i.e., the average of the squared forecast revision when the forecast horizon decreases from h+1 to h. Forecast errors and forecast revisions are based on consensus forecasts for the years 2016 to 2020.

More recently, Bordalo et al. (2020), Kohlhas and Walther (2021), and Angeletos et al. (2021), among others, have provided further evidence for the underreaction of the consensus forecast.

Several studies have found that at the individual level, forecasters tend to overreact to new information. Using data from several countries, Lahiri and Sheng (2008) test for forecaster rationality at the individual level. For each forecast horizon, they consider a regression of individual forecast errors on the current individual forecast revision, whereby the coefficient in front of the revision is the same for all forecasters. For output growth, they find that individual forecasters *overreact* to new information at long forecast horizons but *underreact* at middle-forecast horizons. This result is confirmed in Lahiri and Sheng (2010a) for output growth, while there is only evidence for overreaction at short forecast horizons for inflation.

Bordalo et al. (2020) use SPF and Blue Chip data for several macroeconomic variables and apply the same regression as in Lahiri and Sheng (2008). They also provide evidence for overreaction to news but mainly focus on a forecast horizon of three-quarters-ahead. Importantly, Bordalo et

al. (2020) also run regressions with a forecaster-specific slope parameter. Again, the evidence (as measured by the median parameter estimate across all forecasters) is in line with overreaction at the individual level, which is inconsistent with sticky and noisy information models. However, it is consistent with their model of diagnostic expectations, which can generate underreaction at the consensus level simultaneously with overreaction at the individual level. If individual forecasters have different degrees of overreaction, this is another source of disagreement.

We re-estimated those regressions (with a common slope coefficient for all forecasters) using data from the ZEW survey. In line with Lahiri and Sheng (2008), we find evidence for significant overreaction at most forecast horizons.[6] The evidence provided by Lahiri and Zhao (2020, 3) is more mixed: "Even though overreaction to news at each horizon is more common, the overreaction to individual news is not as ubiquitous as suggested in Bordalo et al. (2020)."

The literature on over-/underreaction in the predictions of professional forecasters is rapidly evolving. For example, Kohlhas and Walther (2021) argue that forecasters simultaneously extrapolate (i.e., overreact to recent realizations of the forecasted variable) and underreact in the sense of a positive correlation between forecast errors and average revisions, which is inconsistent with the diagnostic expectations model of Bordalo et al. (2020). For further evidence, see Kučinskas and Peters (2022) and Juodis and Kučinskas (2023).

Discussion and Conclusion

Various mechanisms can generate heterogeneity in macroeconomic expectations. Some mechanisms can generate disagreement while preserving rational expectations and full

---

[6] Detailed estimation results are omitted for brevity but available upon request.

information, others preserve rational expectations and deviate from full information, while some deviate from both. Empirical evidence from survey data is crucial because it helps to decide which mechanisms are most credible. As Manski (2017, 413) put it: "I favorably observe the increasing willingness of macroeconomic theorists to pose and study alternatives to the RE [rational expectations] assumption, but I worry that models of expectations formation will proliferate endlessly in the absence of empirical research to discipline thinking." Our survey highlights the essential features that drive expectations heterogeneity in observational data, such as prior beliefs or subjective views of the economy. This suggests that more research is needed on the determinants of those beliefs. One aspect that has received less attention but may be essential is the demand side. As Nordhaus (1987) and Lahiri (2012) emphasized, considering the demands of clients or institutional requirements is important for understanding the forecaster's incentives and learning about their true loss functions. Following this line of argument, Valchev and Gemmi (2023) suggest that survey forecasts may not reflect the true expectations of forecasters but might be driven by strategic incentives. Concerning survey design, Patton (2020) argues to communicate explicitly to the forecasters according to which loss function their forecasts will be evaluated. While the literature almost exclusively focuses on disagreement in point predictions, disagreement on forecast uncertainty is an interesting area for future research (see Glas and Hartmann 2022).

The macroeconomic literature has recently focused on regressions that relate forecast errors to forecast revisions. However, as Nordhaus (1987) emphasized, the uncorrelatedness of forecast errors and revisions is only one aspect of a rational forecast. Even if forecasts satisfy this condition, it does not mean that forecasts are of good quality. Conversely, forecasts which fail the rationality test can be highly accurate. For example, we present evidence for underreaction of the consensus forecast. Nevertheless, empirically, the SPF consensus forecasts for many

variables tend to outperform predictions from state-of-the-art econometric models (see, for example, Ang et al. 2007; Faust and Wright 2013). Again, it would be worth putting more effort into understanding what steps in the actual process of forecasting lead to the observed deviations from rationality. For example, parameter estimation uncertainty is typically neglected. In the real world, forecasters must estimate model parameters. In small samples, it can be beneficial in terms of out-of-sample forecast performance to deliberately use a small, misspecified model instead of the true but more complex model. Alternatively, in the presence of non-stationarities and uncertainty about the true model, forecasters will ensure against misspecification by relying on a combined forecast that integrates various models (Bates and Granger 1969; Timmermann 2006). The benefits of forecast combination are illustrated by the fact that the consensus forecast typically outperforms individual forecasts. On top of that, forecasters regularly adjust model-based predictions based on subjective judgment (Brenna and Budrys 2024). Finally, as noted, for example, in Davies and Lahiri (1995), forecasts are contaminated with measurement error. An interesting paper that explores the importance of measurement error and consequences for tests of forecast rationality is Juodis and Kučinskas (2023). All these factors also contribute to explaining disagreement in survey expectations.

References

Afrouzi, Hassan, Spencer Y. Kwon, Augustin Landier, Yueran Ma, and David Thesmar. 2023. "Overreaction in Expectations: Evidence and Theory." *The Quarterly Journal of Economics* 138, no. 3: 1713–1764.

Andrade, Philippe, and Hervé Le Bihan. 2013. "Inattentive Professional Forecasters." *Journal of Monetary Economics* 60, no. 8: 967–982.

Andrade, Philippe, Richard K. Crump, Stefano Eusepi, and Emanuel Moench. 2016. "Fundamental Disagreement." *Journal of Monetary Economics* 83: 106–128.

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. 2022. "Subjective Models of the Macroeconomy: Evidence from Experts and Representative Samples." *The Review of Economic Studies* 89, no. 6: 2958–2991.

Ang, Andrew, Geert Bekaert, and Min Wie. 2007. "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" *Journal of Monetary Economics* 54, no. 4: 1163–1212.

Angeletos, George-Marios, Zhen Huo, and Karthik A Sastry. 2021. "Imperfect Macroeconomic Expectations: Evidence and Theory." *NBER Macroeconomics Annual* 35, no. 1: 1–86.

Batchelor, Roy. 2007. "Bias in Macroeconomic Forecasts." *International Journal of Forecasting* 23, no. 2: 189–203.

Bates, John M., and Clive W. J. Granger. 1969. "The Combination of Forecasts." *Operational Research* 20, no. 4: 451–468.

Benchimol, Jonathan, Makram El-Shagi, and Yossi Saadon. 2022. "Do Expert Experience and Characteristics Affect Inflation Forecasts?" *Journal of Economic Behavior & Organization* 201: 205–226.

Boero, Gianna, Jeremy Smith, and Kenneth F. Wallis. 2015. "The Measurement and Characteristics of Professional Forecasters' Uncertainty." *Journal of Applied Econometrics* 30, no. 7: 1029–1046.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review* 110, no. 9: 2748–2782.

Brenna, Federica, and Zymantas Budrys. 2024. "The Term Structure of Judgement: Interpreting Survey Disagreement." Bank of Lithuania Working Paper Series 123, Bank of Lithuania. URL: https://ideas.repec.org/p/lie/wpaper/123.html.

Brückbauer, Frank, and Michael Schröder. 2022. "The ZEW Financial Market Survey Panel." *Jahrbücher für Nationalökonomie und Statistik* 243, no. 3–4: 451–469.

Capistrán, Carlos. 2008. "Bias in Federal Reserve Inflation Forecasts: Is the Federal Reserve Irrational or Just Cautious?" *Journal of Monetary Economics* 55, no. 8: 1415–1427.

Capistrán, Carlos, and Allan Timmermann. 2009. "Disagreement and Biases in Inflation Expectations." *Journal of Money, Credit and Banking* 41, no. 2–3: 365–396.

Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics* 118, no. 1: 269–298.

Clements, Michael P. 2019. *Macroeconomic Survey Expectations*. Cham: Palgrave Macmillan.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120, no. 1: 116–159.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105, no. 8: 2644–2678.

Conrad, Christian, and Matthias Hartmann. 2019. "On the Determinants of Long-Run Inflation Uncertainty: Evidence from a Panel of 17 Developed Economies." *European Journal of Political Economy* 56: 233–250.

Conrad, Christian, and Matthias Hartmann. 2023. "When Does Information on Forecast Variance Improve the Performance of a Combined Forecast?" Working Paper, Heidelberg University. URL: http://hdl.handle.net/10419/168200.

Conrad, Christian, Zeno Enders, and Alexander Glas. 2022. "The Role of Information and Experience for Households' Inflation Expectations." *European Economic Review* 143: 104015.

Conrad, Christian, Julius T. Schoelkopf, and Nikoleta Tushteva. 2023. "Long-Term Volatility Shapes the Stock Market's Sensitivity to Macroeconomic News." DOI: 10.2139/ssrn.4632733.

Christian Conrad, Julius T. Schoelkopf, Michael Weber, and Frank Brückbauer. Forthcoming 2024. "Beyond the Numbers: How Narratives about Inflation Shape Professional Forecasters' Stock Market Expectations." Working Paper, Heidelberg University. Davies, Anthony, and Kajal Lahiri. 1995. "A New Framework for Analyzing Survey Forecasts using Three-Dimensional Panel Data." *Journal of Econometrics* 68, no. 1: 205–227.

Diebold, Francis X., and Jose A. Lopez. 1996. "Forecast Evaluation and Combination." In *Handbook of Statistics*, edited by G. S. Maddala and C. R. Rao, 14: 241–268. Amsterdam: Elsevier.

Dovern, Jonas. 2015. "A Multivariate Analysis of Forecast Disagreement: Confronting Models of Disagreement with Survey Data." *European Economic Review* 80: 16–35.

Dovern, Jonas, and Matthias Hartmann. 2017. "Forecast Performance, Disagreement, and Heterogeneous Signal-to-Noise Ratios." *Empirical Economics* 53, no. 1: 63–77.

Dovern, Jonas, and Johannes Weisser. 2011. "Accuracy, Unbiasedness and Efficiency of Professional Macroeconomic Forecasts: An Empirical Comparison for the G7.*" International Journal of Forecasting* 27, no. 2: 452–465.

Dovern, Jonas, Ulrich Fritsche, and Jiri Slacalek. 2012. "Disagreement among Forecasters in G7 Countries." *Review of Economics and Statistics* 94, no. 4: 1081–1096.

Elliott, Graham, Allan Timmermann, and Ivana Komunjer. 2005. "Estimation and Testing of Forecast Rationality under Flexible Loss." *Review of Economic Studies* 72, no. 4: 1107–1125.

Elliott, Graham, Ivana Komunjer, and Allan Timmermann. 2008. "Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss?" *Journal of the European Economic Association* 6, no. 1: 122–157.

Engelberg, Joseph, Charles F. Manski, and Jared Williams. 2009. "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters." *Journal of Business & Economic Statistics* 27, no. 1: 30–41.

Faust, Jon, and Jonathan H. Wright. 2013. "Forecasting Inflation." In *Handbook of Economic Forecasting*, edited by G. Elliott and A. Timmermann, 2: 2–56. Amsterdam: Elsevier.

Gemmi, Luca, and Rosen Valchey. 2023. "Biased Surveys." NBER Working Paper Series no. w31607. SSRN: https://ssrn.com/abstract=4553611.

Giacomini, Raffaella, Vasiliki Skreta, and Javier Turen. 2020. "Heterogeneity, Inattention, and Bayesian Updates." *American Economic Journal: Macroeconomics*, 12, no. 1: 282–309.

Giordani, Paolo, and Paul Söderlind. 2003. "Inflation Forecast Uncertainty." *European Economic Review* 47, no. 6: 1037–1059.

Glas, Alexander. 2020. "Five Dimensions of the Uncertainty-Disagreement Linkage." *International Journal of Forecasting* 36, no. 2: 607–627.

Glas, Alexander, and Matthias Hartmann. 2022. "Uncertainty Measures from Partially Rounded Probabilistic Forecast Surveys." *Quantitative Economics* 13, no. 3: 979–1022.

Granger, Clive W. J. 1969. "Prediction with a Generalized Cost of Error Function." *Operational Research* 20: 199–207.

Hur, Joonyoung, and Insu Kim. 2016. "Information Rigidity in Survey Data: Evidence from Dispersion in Forecasts and Forecast Revisions." *Economic Letters* 142: 1–14.

Isiklar, Gultekin, and Kajal Lahiri. 2007. "How Far Ahead Can We Forecast? Evidence from Cross-Country Surveys." *International Journal of Forecasting* 23, no. 2: 167–187.

Juodis, Artūras, and Simas Kučinskas. 2023. "Quantifying Noise in Survey Expectations." *Quantitative Economics* 14, no. 2: 609–650.

Kohlhas, Alexandre N., and Ansgar Walther. 2021. "Asymmetric Attention." *American Economic Review* 111, no. 9: 2879–2925.

Krüger, Jens J., and Julian LeCrone. 2019. "Forecast Evaluation under Asymmetric Loss: A Monte Carlo Analysis of the EKT Method." *Oxford Bulletin of Economics and Statistics* 81, no. 2: 437–455.

Kučinskas, Simas, and Florian S. Peters. 2022. "Measuring Under- and Overreaction in Expectation Formation." *The Review of Economics and Statistics*, forthcoming. SSRN: https://ssrn.com/abstract=3188065.

Lahiri, Kajal. 2012. "Comment on 'Rationality Tests Based on Multi-Horizon Bounds'." *Journal of Business & Economic Statistics* 30, no. 1: 20–25.

Lahiri, Kajal, and Xuguang Sheng. 2008. "Evolution of Forecast Disagreement in a Bayesian Learning Model." *Journal of Econometrics* 144, no. 2: 325–340.

Lahiri, Kajal, and Xuguang Sheng. 2010a. "Learning and Heterogeneity in GDP and Inflation Forecasts." *International Journal of Forecasting* 26, no. 2: 265–292.

Lahiri, Kajal, and Xuguang Sheng. 2010b. "Measuring Forecast Uncertainty by Disagreement: The Missing Link." *Journal of Applied Econometrics* 25, no. 4: 514–538.

Lahiri, Kajal, and Yongchen Zhao. 2020. "The Nordhaus Test with Many Zeros." *Economic Letters* 193: 109308.

Lucas, Robert E. Jr. 1976. "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy* 1: 19–46.

Lucas, Robert E. Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4, no. 2: 103–124.

Lucas, Robert E. Jr. 1973. "Some International Evidence on Inflation-Output Tradeoffs." *American Economic Review* 63, no. 3: 326–334.

Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Quarterly Journal of Economics* 126, no. 1: 373–416.

Malmendier, Ulrike, and Stefan Nagel. 2016. "Learning from Inflation Experiences." *Quarterly Journal of Economics* 131, no. 1: 53–87.

Malmendier, Ulrike, Stefan Nagel, and Zhen Yan. 2021. "The Making of Hawks and Doves." *Journal of Monetary Economics* 117: 19–42.

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117, no. 4: 1295–1328.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2003. "Disagreement about Inflation Expectations." *NBER Macroeconomics Annual* 18: 209–248.

Manski, Charles F. 2017. "Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise." *NBER Macroeconomics Annual* 32: 411–471.

Mincer, Jacob A., and Victor Zarnowitz. 1969. "The Evaluation of Economic Forecasts." In *Economic Forecasts and Expectations*, 3–46. New York, NY: National Bureau of Economic Research.

Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica* 29, no. 3: 315–335.

Nordhaus, William D. 1987. "Forecasting Efficiency: Concepts and Applications." *Review of Economics and Statistics* 69, no. 4: 667–674.

Patton, Andrew J. 2020. "Comparing Possibly Misspecified Forecasts." *Journal of Business & Economic Statistics* 38, no. 4: 796–809.

Patton, Andrew J., and Allan Timmermann. 2007. "Properties of Optimal Forecasts under Asymmetric Loss and Nonlinearity." *Journal of Econometrics* 140, no. 2: 884–918.

Patton, Andrew J., and Allan Timmermann. 2010. "Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-Sectional Dispersion." *Journal of Monetary Economics* 57, no. 7: 803–820.

Pesaran, M. Hashem, and Martin Weale. 2006. "Survey Expectations." In *Handbook of Economic Forecasting*, edited by G. Elliott, C. W. J. Granger and A. Timmermann, 1: 715–776. Amsterdam: Elsevier.

Rich, Robert, and Joseph Tracy. 2021. "A Closer Look at the Behavior of Uncertainty and Disagreement: Micro Evidence from the Euro Area." *Journal of Money, Credit and Banking* 53, no. 1: 233–253.

Sheng, Xuguang Simon. 2015. "Evaluating the Economic Forecasts of FOMC Members." *International Journal of Forecasting* 31, no. 1: 165–175.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50, no. 3: 665–690.

Theil, Henri. 1966. *Applied Economic Forecasting*. Amsterdam: North-Holland.

Timmermann, Allan. 2006. "Chapter 4: Forecast Combinations." In *Handbook of Economic Forecasting*, edited by G. Elliott, C. W. J. Granger and A. Timmermann, 1: 135–196. Amsterdam: Elsevier.

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion. 2022. "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications." *Journal of Economic Perspectives* 36, no. 3: 157–184.

Zarnowitz, Victor. 1985. "Rational Expectations and Macroeconomic Forecasts." *Journal of Business & Economic Statistics* 3, no. 4: 293–311.

Zarnowitz, Victor, and Louis A. Lambros. 1987. "Consensus and Uncertainty in Economic Predictions." *Journal of Political Economy* 95, no. 3: 591–621.

Zellner, Arnold. 1986. "Biased Predictors, Rationality and the Evaluation of Forecasts." *Economics Letters* 21, no. 1: 45–48.

Zellner, Arnold. 2002. "Information Processing and Bayesian Analysis." *Journal of Econometrics* 107, no. 1–2: 41–50.