FORECASTER (MIS-)BEHAVIOR

Tobias Broer and Alexandre N. Kohlhas*

Abstract—We document two stylized facts in expectational survey data. First, professional forecasters overrevise their macroeconomic expectations. Second, such overrevisions mask evidence of both over- and underreactions to public signals. We show that the first fact is inconsistent with standard models of noisy rational expectations, but consistent with behavioral and strategic models. The second fact, in contrast, presents a puzzle for existing theories. We propose an extension of noisy rational expectations that allows forecasters to be overconfident in their information. We show that this feature when combined with the endogeneity of public signals leads to over- and undereactions consistent with the data.

I. Introduction

EXPECTATIONS are central to economics. Because individual expectations are typically unobserved, however, it is often difficult to discriminate between alternative models of expectation formation. One exception is professional forecaster surveys, which regularly publish individual expectations about macroeconomic and financial variables. Indeed, Muth (1961) proposed the rational-expectation theory in part to explain the perceived sluggishness of average survey expectations as a rational response to noisy information (Muth, 1961, p. 316).

Although the full-information variant of rational expectations later became the benchmark of modern macroeconomics, the work of Woodford (2002), Sims (2003), and others,[1] has revived interest in noisy information models of rational expectations. In turn, this has rekindled interest in the use of survey data to better discipline and test such models. In line with a central prediction of noisy rational expectations, Coibion and Gorodnichenko (2012, 2015) recently document that the average of survey expectations underreacts to new information relative to what a full-information framework would prescribe.[2] However, such underreactions of *average* expectations are not only consistent with noisy rational expectations, but also with a host of other both rational and behavioral theories of *individual* expectation formation.

In this paper, we provide new evidence on the statistical properties of individual survey expectations of macroeconomic variables. We document two stylized facts that present a challenge for noisy rational expectations. First, individual forecasters *overrevise* their macroeconomic expectations. Second, such overrevisions mask both *over- and underreactions* to salient public signals. We show that the first fact is inconsistent with standard models of noisy rational expectations, but in line with, for example, models of strategic forecaster behavior. The second fact, in contrast, presents a puzzle for existing theories of expectation formation.

Our main contribution is to explain this coexistence of over- and underreactions: We propose a simple extension of noisy rational expectations that allows forecasters to be overconfident in the precision of their own information (both relative to the truth and relative to their perception of others). We show that such overconfidence causes forecasters to overrevise their expectations and misperceive others' responses to information. Importantly, such misperception leads forecasters to misinterpret public signals that aggregate others' information and actions, and results in over- or underreactions that are consistent with the data.

A well-known consequence of rational (mean-squared optimal) expectations is that individual forecast errors should be unpredictable based on known information. The two stylized facts that motivate our theory test this prediction. Our first test relates individual forecast errors to individual revisions in fixed-date forecasts.[3] Basic introspection by rational forecasters requires these two to be uncorrelated, even in the presence of noisy information. Our second test instead exploits the survey data to relate individual forecast errors directly to elements of public information that are salient to professional forecasters. Once more, individual rationality requires individual forecast errors to be uncorrelated with these.

As a benchmark, we first consider inflation forecasts from the Survey of Professional Forecasters (SPF). We use the outcomes of our tests to document two stylized facts.

First, individual forecast revisions are systematically too large. This manifests itself in a pronounced negative relationship between individual forecast errors, on the one hand, and individual forecast revisions, on the other hand. Second, such observed overrevisions mask evidence of both over- and underreactions to salient public signals that are both predictive about future inflation, relevant, and observed in real time (e.g., previous consensus forecasts or changes in

Received for publication April 28, 2021. Revision accepted for publication March 30, 2022. Editor: Olivier Coibion.

*Broer: Paris School of Economics; Kohlhas: University of Oxford.

We are grateful to the editor, Olivier Coibion, and an anonymous referee for extensive feedback that improved the paper. We also thank George-Marios Angeletos, Vladimir Asriyan, Harris Dellas, Zeno Enders, Zhen Huo, Per Krusell, Kurt Mitman, Kristoffer Nimark, Torsten Persson, Andrei Shleifer, Peter Sørensen, Jonathan de Quidt, and seminar and conference participants at the Barcelona 2019 Summer Forum, Bern University, CREi, Danish Central Bank, ECB, ESSIM 2020, Heidelberg University, Oxford University, Oslo University, PSE, Riksbank, and SED 2019 for their comments. Financial support from *Ragnar Söderbergs Stiftelsen* is gratefully acknowledged.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01210.

[1]See, for example, Mankiw and Reis (2002), Angeletos and Pavan (2007), Lorenzoni (2009), and Maćkowiak and Wiederholt (2009).

[2]Coibion and Gorodnichenko (2012), Andrade and Le Bihan (2013), Bordalo et al. (2020), and Kohlhas and Walther (2021) document related evidence.

[3]In contemporaneous and independent work, Bordalo et al. (2020) propose a similar test. The working paper version of Bordalo et al. (2020) (Bordalo et al., 2018b, p. 7) acknowledges this simultaneity. We discuss the similarities and differences between the two approaches in the related literature section.

The Review of Economics and Statistics, September 2024, 106(5): 1334–1351

https://doi.org/10.1162/rest_a_01210

the unemployment rate). We document that these patterns extend to forecasts of macroeconomic variables other than inflation, different forecast horizons, to the euro area, as well as to other forecasters than those that label themselves professional.

Combined, our empirical results present a challenge for existing models of expectation formation. While simple models of noisy rational expectations are consistent with an underrevision of average expectations, they are *prima facie* inconsistent with the overrevision documented at the individual level. Alternative theories of forecaster behavior that incorporate (i) the specific strategic considerations faced by professional forecasters (e.g., Laster et al., 1999; Ottaviani & Sørensen, 2006; Ehrbeck & Waldmann, 1996); (ii) common behavioral biases (e.g., Daniel et al., 1998; Bordalo et al., 2020); or (iii) trembling-hand noise can explain such overrevisions. However, as we show, these theories either all predict optimal use of public information (conditional on private information), or that forecasters overreact to *all* new information, irrespective of its source. Both predictions are inconsistent with the simultaneous over- and underreactions to public information that we document in the data.[4]

To account for our empirical results, we propose a simple extension of noisy rational expectations that allows forecasters to be overconfident in their own information. Specifically, we allow forecasters to both perceive their private information to be more precise than it actually is ("*absolute overconfidence*"; Alpert & Raiffa, 1982; Soll & Klayman, 2004; and others), and to be more precise than the information of others ("*relative overconfidence*"; Alicke & Govorun, 2005; Larrick et al., 2007). Both dimensions of overconfidence are commonly used in the psychology literature (Moore & Healy, 2008), and we find direct evidence for them in the survey data. We show that, taken together, absolute and relative overconfidence can explain our empirical results when combined with the central feature that most public signals reflect the aggregate outcome of others' choices, and hence their information.

All else equal, absolute overconfidence makes forecasters overreact to private information, and hence makes individual forecast revisions too large. In contrast, relative overconfidence makes forecasters underestimate the precision of others' information. This is important.

Any equilibrium model of expectation formation that internalizes that most public signals reflect others' choices requires an assumption about people's views about others' information. Rational expectations solves this problem by imposing the "symmetry assumption": that others' information is equal in quality to an individual agent's. Relative overconfidence, by contrast, imposes the empirically motivated "better than others" perception. By underestimating the precision of others' information, relative overconfidence causes forecasters to expect public signals to respond less to new information and to be less precise.

We show that such misperceptions have two offsetting effects: Underestimating the precision of public signals, all else equal, leads forecasters to dismiss them, and underreact to their realizations. However, underestimating the responsiveness of public signals, by contrast, leads forecasters to over-infer information from any given signal realization and hence to overreact.

We demonstrate these results within the context of a workhorse noisy information model with mean-squared error preferences (e.g., Veldkamp, 2011). Although our model is simple, we quantitatively validate it along three dimensions.

First, we show that our model can match the estimated overrevision of inflation forecasts at the same time as the estimated overreaction to a particular public signal, previous period's consensus forecast. We focus on consensus forecasts because it reflects a public signal that only aggregates others' information. This allows us to focus on overconfidence's role in creating a friction between forecasters' perception of a public signal and that which arises in equilibrium. As argued in Ottaviani and Sørensen (2006), consensus forecasts also represent a particularly salient public signal for professional forecasters, such as those in the SPF. Second, an attractive feature of the survey data on professional forecasters is that respondents also report forecast densities, in addition to point estimates. We show that this additional information, when combined with auxiliary data on higher-order expectations from Coibion et al. (2021), allows us to validate our assumptions of absolute and relative overconfidence in the survey data. Third, two key implications of our model are that (i) forecasters should underreact more to public signals that are less precise, and (ii) that the magnitude of over- and underreactions should change with the volatility of the forecasted variable. We demonstrate that both predictions are in line with the patterns of responses in the data. We conclude the paper by studying the implications of our model for the distribution of forecast errors.

Finally, professional forecasters may admittedly differ from other economic agents in their incentives and information about the state of the economy. In this paper, we confront this issue head-on by directly contrasting the ability of agency-based models to explain the observed under- and overreactions with simple behavioral alternatives. To the extent that the evidence we uncover below speaks in favor of widely documented behavioral biases, rather than particular strategic incentives, we think that our results should carry over to other contexts. Indeed, we provide some evidence to this effect later in the paper.

A. *Related Literature*

Our paper is related to several strands of research. We review these in order of proximity, starting with the most closely related.

[4]Angeletos and Huo (2021) show that the overrevisions that we document are also inconsistent with two other alternatives: "cognitive discounting" and "level-k thinking."

First, our paper relates to studies that use expectational survey data to test models of noisy rational expectations.[5] Recently, Coibion and Gorodnichenko (2012, 2015) document that *average forecasts* of several macroeconomic variables, across different surveys, underreact to new information. Our study departs from this observation, and studies both average and individual-level forecasts within a unified framework. Complementary to our paper, in contemporaneous and independent work, Bordalo et al. (2020) demonstrate similar overrevisions of *individual-level* forecasts to those that we document. In contrast to their paper, we show that these overrevisions mask evidence of both over- and underreactions to salient public signals. We further show that such simultaneous over- and underreactions present a challenge for existing models, including Bordalo et al.'s (2020) theory of "diagnostic expectations." Closely related, Kohlhas and Walther (2021) also depart from Coibion and Gorodnichenko's (2015) observation, but demonstrate that individual forecasters simultaneously extrapolate from recent events. Such extrapolation can be viewed as an overreaction to a specific public signal: that of the past outcome of the forecasted variable. Our paper elaborates on this observation and studies forecasters' responses to a wide set of salient public signals. Importantly, we demonstrate that forecasters also occasionally *underreact* to publicly available information. Finally, building on the above work and our present study, Angeletos et al. (2021) propose a model that combines absolute overconfidence with an additional behavior friction (overpersistence). This model tractably speaks to the above evidence, as well as several auxiliary results.[6] We view the above strand of literature as presenting complementary and related steps towards a unified model of expectations that is consistent with the survey data.

Second, although forecaster information is sometimes acknowledged to be an upper bound of that held by the population at large (Marinovic et al., 2013), most studies abstract from the particular characteristics that separate professional forecasters from the rest of the population. This has attracted criticism (e.g., Scharfstein & Stein, 1990 and Lamont, 2002) and given rise to a literature that looks at forecasters' incentives to distort their stated predictions (e.g., Laster et al., 1999; Ehrbeck & Waldmann, 1996). Our contribution in this context is to show, within a common framework, how several of the most prominent of such agency-based models are inconsistent with individual-level forecasts from a variety of professional surveys.

Third, our paper relates to the substantial body of work that links over- and underreaction of expectations to asset price anomalies. For example, Daniel et al. (1998) show how a model of overprecision (leading to overreactions) and self-attribution of skill (leading to underreactions) is consistent with the excess volatility and short-run momentum often found in financial markets. Barberis et al. (1998), in contrast, show how a model of conservatism (underreaction) and "representativeness" (overreaction) can explain the underreaction of stock prices to earnings announcements jointly with the overreaction of stock prices to extreme events. Lastly, and closely related to our notion of relative overconfidence, Eyster et al. (2019) show how "cursedness" (the failure to extract information from market prices) may explain momentum in asset prices. In contrast to this work, our evidence of over- and underreactions is based directly on forecasters'stated predictions rather than the behavior of equilibrium objects, such as asset prices. We hence view our evidence as a useful anchor for these models.

II. Empirical Evidence

In this section, we document three features of U.S. inflation forecasts. We show that professional forecasters' average inflation expectations underreact to information received between two periods. We then show that, at the individual level, the same forecasters by contrast make forecast revisions that are too large. Lastly, we document that the overrevisions at the individual level mask evidence of both over- and underreactions to salient public signals.

A. Data

We focus on U.S. inflation forecasts from the *Survey of Professional Forecasters* (SPF).[7] At the start of each quarter, the SPF asks its respondents for their forecasts of a number of key macroeconomic and financial variables, and publishes them, in anonymous format but with personal identifiers, shortly thereafter. We study SPF forecasts of the year-on-year percentage change in the GNP/GDP deflator, for which the survey includes consistent forecasts for the six quarters following the survey quarter. We focus on inflation forecasts for three reasons. First, because inflation expectations play a central role in the economy as determinants of wages, goods and asset prices. Second, to compare our estimates to those of previous studies, which have focused disproportionally on inflation. And third, because data on inflation forecasts are available for a substantially longer time-span than forecasts of other variables. Throughout, we consider first-release realizations of inflation to most accurately capture the precise definition of the variable being forecasted. Importantly for our purposes, although the precise schedule over the quarter has changed over time, the administrators of the SPF have

[5]Apart from the implications discussed here, broad aspects of survey forecasts are clearly consistent with noisy rational expectations. First, survey forecasts are dispersed. Second, forecasts are often smoother than the variable that is being forecasted. In fact, one of Muth's (1961) aims in proposing the rational expectations hypothesis was to explain these two stylized facts (p. 316 in Muth, 1961).

[6]Angeletos and Huo (2021) show that the approach proposed in this paper also has the advantage that the *as-if* myopia and anchoring that are consequences of noisy rational expectations directly extend to our model of overconfidence.

[7]The SPF is the oldest quarterly survey of individual macroeconomic forecasts in the United States, dating back to 1968 (Croushore, 1993).

consistently and publicly published the average of survey results well before sending out the next round of the questionnaire.[8] The information set of respondents therefore includes the consensus (or average) forecast from the previous quarter.

B. Average Forecasts

We first study the properties of average inflation forecasts. We denote respondent i's forecast made in period t of inflation π in period $t+h$ as $f_{\mathrm{it}}\pi_{t+h}$. We then calculate the average forecast as $f_t\pi_{t+h} = \frac{1}{N_t}\sum_{i=1,\ldots,N_t} f_{\mathrm{it}}\pi_{t+h}$, where N_t denotes the number of respondents in period t. A respondent's forecast error is $\pi_{t+h} - f_{\mathrm{it}}\pi_{t+h}$, while the average forecast error is $\pi_{t+h} - f_t\pi_{t+h}$.

A well-known implication of full information and rational expectations (with mean-squared error preferences) is that forecast errors should be uncorrelated with known information. Our first test explores this prediction by estimating the Coibion and Gorodnichenko (2015) regression:

$$\pi_{t+h} - f_t\pi_{t+h} = a + b(f_t\pi_{t+h} - f_{t-1}\pi_{t+h}) + v_t, \qquad (1)$$

where $f_t\pi_{t+h} - f_{t-1}\pi_{t+h}$ denotes the average forecast revision between period $t-1$ and t, and v_t an error term. We hence estimate the relationship between average errors and average revisions. Panel b in figure 1 presents the results for one-year ahead inflation forecasts ($h = 4$).

Average revisions are positively correlated with average errors ($b > 0$). This effect is strong and highly significant, in line with the results in Coibion and Gorodnichenko (2015), among others. On average, forecasters underrevise their expectations relative to the full information rational expectations benchmark, leading to a positive correlation between average errors, on the one hand, and average revisions, on the other hand.

Although inconsistent with full (common) information and rational expectations, $b > 0$ is in line with several popular models of rational expectations that allow for individual-specific noise in respondents' information (Coibion & Gorodnichenko, 2015).[9] In such noisy information models, forecasters revise their expectations by less than a hypothetical agent would do if she could observe the average information in the population. This is because forecasters rationally downweigh their own information to account for its noisiness. However, because individual-specific noise terms cancel on average, this downweighing of new information leads to a positive correlation between average errors, on the one hand, and average revisions, on the other hand ($b > 0$). Relative to the full-information benchmark, noisy private information leads to an underrevision of average forecasts in response to average information.

We next turn to the statistical properties of individual inflation forecasts. We continue to explore the implication that rational errors should be orthogonal to known information.

C. Individual Forecasts

Overrevision of individual forecasts. Equation (1) studies the relationship between errors and *average revisions*.[10] However, basic introspection implies that *individual revisions* are always known to individual forecasters. This is even in the presence of noisy private information. This, in turn, implies that even with noisy information, rational (mean-squared optimal) errors should be uncorrelated with individual revisions. Estimates of the slope coefficient in equation (1) at the individual level should equal zero.

Panel a in figure 1 shows that this implication is *prima facie* not borne out by the data. The conditional means of individual errors are negatively associated with the means of individual revisions (left panel), suggesting a negative relationship. To test this implication, we estimate a version of equation (1) at the individual level, using the benchmark specification:

$$\pi_{t+h} - f_{\mathrm{it}}\pi_{t+h} = \alpha_i + \beta\,(f_{\mathrm{it}}\pi_{t+h} - f_{it-1}\pi_{t+h}) + \nu_{\mathrm{it}}, \qquad (2)$$

where α_i denotes a respondent fixed effect. Panel b in figure 1 confirms our initial impressions.

The estimate of β is significantly negative and numerically large, inconsistent with the predictions of (noisy) rational expectations. This negative estimated value of β implies that positive individual revisions are associated with negative errors. Forecasters on average revise their forecasts by too much relative to the rational expectations benchmark, and hence on average overreact to the information received between subsequent survey rounds.

However, importantly, such overrevisions of individual forecasts do not inform us about the composition of responses that lead to a negative estimate for β. All we can conclude is that forecasters overreact *on average*. In particular, estimates of equation (2) do not allow us to separate between (i) whether the overrevision of expectations is comprised exclusively of overreactions to new information, or (ii) whether the overall overrevision masks evidence of both over- and underreactions. As we argue in section III, this distinction is important for our analysis, as it will greatly constrain the set of models that are consistent with the data.

Finally, notice that a positive estimate of b in figure 1 corresponds to an average underrevision relative to the *full information and rational expectations* benchmark. By

[8] See p. 8 in the documentation: https://www.philadelphiafed.org/-/media/research-and-data/real-time-center/survey-of-professional-forecasters/.

[9] We note that with noisy *public information*, least-squares estimates of b in equation (1) are downwardly biased. As argued in Coibion and Gorodnichenko (2015) online appendix A, such downward bias, however, still entails that statistically significant findings of $b > 0$ imply average underreactions relative to full information.

[10] We further note that equation (1) is equivalent to estimating the linear relationship between *individual* errors and *average* revisions. This is because $\mathbb{C}\mathrm{ov}(\pi_{t+h} - f_t\pi_{t+h}, x_t) = \frac{1}{N}\sum_{i=1}^{N}\mathbb{C}\mathrm{ov}(\pi_{t+h} - f_{\mathrm{it}}\pi_{t+h}, x_t)di$ for some common variable x_t and assuming N *ex-ante* identical forecasters. In equation (1), $x_t \equiv f_t\pi_{t+h} - f_{t-1}\pi_{t+h}$.

FIGURE 1.—OVER- AND UNDERREACTIONS IN THE SURVEY OF PROFESSIONAL FORECASTERS

Panel a: Conditional Forecast Errors

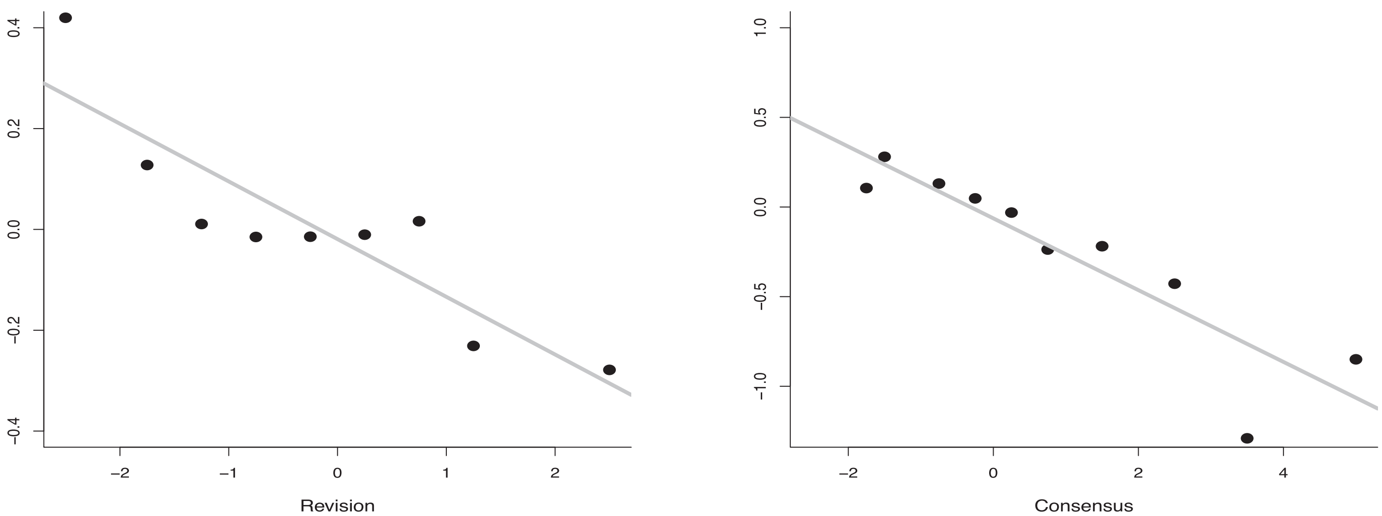

The left-hand side panel depicts (on the vertical axis) the average of individual forecast errors taken within equally-sized bins of the distribution of individual forecast revisions (horizontal axis). The right-hand side panel shows (on the vertical axis) the average of individual forecast errors this time taken within bins of the distribution of consensus forecasts from the previous wave of the SPF (horizontal axis). All variables are demeaned by subtracting their (individual) averages during the SPF sample period (1970Q1–2020Q1).

Panel b: Regression Estimates

	Average forecasts	Individual forecasts		
	Forecast error	Forecast error	Forecast error	Forecast error
Forecast revision	1.118***	−0.199***	–	−0.206***
	(0.287)	(0.067)		(0.067)
Previous consensus	–	–	−0.192**	−0.200**
			(0.085)	(0.088)
Constant	−0.054	–	–	–
	(0.073)			
Observations	196	5,480	5,675	5,480
F Statistic	44.067	113.66	118.98	122.67
R^2	0.185	0.021	0.022	0.045

Estimates of equations (1), (2), and (3) using SPF forecasts of one-year ahead inflation ($h = 4$). Column 1 presents estimates with a constant term. Columns 2–4 include individual (respondent) fixed effects. Robust (double-clustered) standard errors in parentheses. Sample: 1970Q1–2020Q1. $^{*}p < 0.1$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

contrast, a negative estimate of β suggests an overrevision (or overreaction) relative to the *rational expectations* case, allowing for the presence of *noisy information*. We use both notions of over- and underrevisions interchangeably below when there is no cause for confusion.

Over- and underreactions to public signals. In order to provide a first pass at a breakdown of the composition of responses that lead to $\beta < 0$ in figure 1, our third test considers the relationship between errors and the public signals that forecast revisions are based on. (We focus on public signals because those are also observed by researchers.) In particular, we estimate the following regression equation:

$$\pi_{t+h} - f_{\text{it}}\pi_{t+h} = \alpha_i + \delta y_t + \nu_{\text{it}}, \tag{3}$$

where α_i denotes a respondent fixed effect and y_t a public signal that is observed by forecasters. The third implication of rational expectations that we focus on is that δ should equal zero, as any nonzero coefficient would contradict the assumption that public information is used efficiently. The *Law of Iterated Expectations* implies that if the public signal is included in forecasters' information sets $y_t \in \Omega_{\text{it}}$ and forecasts are rational $f_{\text{it}}\pi_{t+h} = \mathbb{E}[\pi_{t+h} \mid \Omega_{\text{it}}]$, then there should always be zero correlation between errors and the public signal.[11]

The predicted zero association between rational errors and public information also allows for a clean interpretation of any nonzero δ−estimates. Because $\pi_{t+k} - \mathbb{E}[\pi_{t+k} \mid \Omega_{\text{it}}]$ is uncorrelated with y_t, we can add and subtract $\mathbb{E}[\pi_{t+k} \mid \Omega_{\text{it}}]$ from the left-hand side of equation (3). This shows that δ is positive (negative) if and only if the rational expectations

[11] In particular, if respondents are rational:

$$\begin{aligned}\alpha_i + \delta \times y_t &= \mathbb{E}[\pi_{t+h} - f_{\text{it}}\pi_{t+h} \mid y_t] \\ &= \mathbb{E}[\pi_{t+h} \mid y_t] - \mathbb{E}\{\mathbb{E}[\pi_{t+h} \mid \Omega_{\text{it}}] \mid y_t\} \\ &= \mathbb{E}[\pi_{t+h} \mid y_t] - \mathbb{E}[\pi_{t+h} \mid y_t] = 0 + 0 \times y_t,\end{aligned} \tag{4}$$

where the third equality follows from $y_t \in \Omega_{\text{it}}$ and the Law of Iterated Expectations. Hence, $\delta = 0$. Notice that one strength of the approach in equation (3) is that it allows forecasters to *rationally* choose to disregard other public signals $z_t \neq y_t$, including components that make up the state of inflation. The only requirement is that y_t is included in the information sets.

FIGURE 2.—INFLATION FORECAST ERRORS AND DIFFERENT PUBLIC SIGNALS

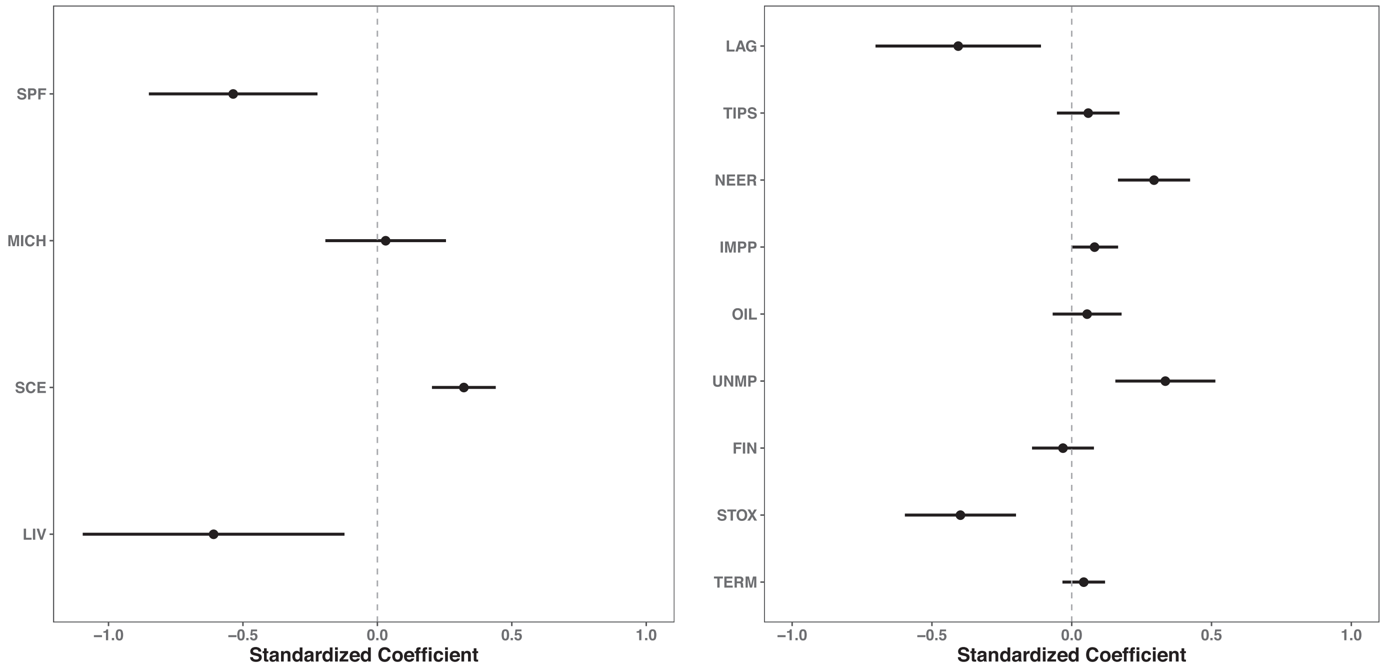

The figure depicts estimates of δ in equation (3) (horizontal axis) for various public signals (vertical axis). The left-hand side panel shows the coefficient estimates for previous period's consensus estimate of one-year ahead inflation ($h = 4$) from the Survey of Professional Forecasters (SPF), the Michigan Survey of Consumers (MICH), the Survey of Consumer Expectations (SCE), and the Livingston Survey (LIV). The right-hand side panel shows estimates of δ using one-period lagged inflation outcomes (LAG), 10-year inflation expectations from the TIPS market (TIPS), the year-over-year change in the nominal effective exchange rate (NEER), the year-over-year change in import prices (IMP), the year-over-year change in the WTI oil price (OIL), the unemployment rate (U), the Cleveland Fed's Financial Market-based measure of future inflation (FIN), the log-linear real detrended level of the SP500 (STOX), and the 10-year-2-year term spread (TERM). All variables have been standardized, and have been signed such that an increase predicts higher inflation one year out. All variables and growth rates have also been derived using the latest available data at the time of the inflation forecast. Whisker-intervals correspond to 95% robust doubled-clustered confidence bounds. Online appendix table B.1 provides further details on the estimates.

forecast $\mathbb{E}[\pi_{t+k} \mid \Omega_{it}]$ has a stronger (weaker) reaction to the public signal y_t than the actual forecast $f_{it}\pi_{t+h}$.[12] Consistent with our earlier use of the terms, we say that forecasters *overreact* to y_t if $\delta < 0$. Conversely, we say that forecasters *underreact* if $\delta > 0$.

To estimate equation (3) requires a particular piece of public information that is at the same time publicly observed, relevant, and salient to professional forecasters. We first focus on a natural example of such public information within our context: that of the consensus forecast from the previous wave of the survey ($y_t = f_{t-1}\pi_{t+h}$). As argued in the introduction, and forcefully in Ottaviani and Sørensen (2006), professional forecasters pay close attention to realizations of consensus. This is to assess how well they perform relative to their competitors. Consensus forecasts should therefore provide a conservative benchmark against which to test the orthogonality of individual forecast errors to public information.

Panel a in figure 1 (right panel) depicts the conditional means of individual forecast errors of one-year ahead inflation ($h = 4$), and shows that these decrease in previous period's consensus forecast. Panel b in figure 1 confirms this impression. The estimate of δ in equation (3) is negative and statistically significant, inconsistent with rational expectations. Individual errors are, on average, more negative not only when individual revisions are more positive, but also when the previous consensus forecast is higher. We conclude that forecasters appear to overreact to the information contained in consensus forecasts. These overreactions are corroborated in the final column of panel b in figure 1, where we report the coefficient estimates from a multivariate regression that includes both individual forecast revisions and consensus. These estimates suggest that even conditional on individual revisions forecasters overreact to consensus.

The negative estimate of δ in figure 1 may suggest that forecasters overreact to all information (implying $\delta < 0$ for all public signals). However, figure 2 shows that such uniformity does not exist. The figure presents estimates of δ from equation (3) using a variety of public signals. We divide this evidence into two types: (i) alternative survey measures of future inflation, similar to consensus forecasts (left-hand side panel) and (ii) other publicly observable time series that are often used to predict inflation (right-hand side panel). We take the latter set of variables from the European Central Bank's published list of "important inflation indicators," to tie our hands with respect to variable selection.[13] A similar set of variables are used in Cecchetti (1995), Canova (2007),

[12]Specifically, we have that

$$\begin{aligned}\alpha_i + \delta \times y_t &= \mathbb{E}[\pi_{t+h} - f_{it}\pi_{t+h} \mid y_t] \\ &= \mathbb{E}[\pi_{t+h} - \mathbb{E}[\pi_{t+k} \mid \Omega_{it}] + \mathbb{E}[\pi_{t+k} \mid \Omega_{it}] - f_{it}\pi_{t+h} \mid y_t] \\ &= \mathbb{E}[\mathbb{E}[\pi_{t+k} \mid \Omega_{it}] - f_{it}\pi_{t+h} \mid y_t].\end{aligned}$$

[13]See, for example, https://www.ecb.europa.eu/pub/pdf/other/ebart2017 04_01.en.pdf. The main difference is that we avoid measures of the "output

TABLE 1.—ROBUSTNESS AND ALTERNATIVE ESTIMATES

	Avg. forecast error		Ind. forecast error			
Description	b-coef	Std. error	β-coef	Std. error	δ-coef	Std. error
GDP deflator (SPF)	1.118	(0.287)	−0.199	(0.067)	−0.192	(0.085)
CPI inflation (SPF)	0.282	(0.230)	−0.293	(0.098)	−0.461	(0.079)
Real GDP (SPF)	0.783	(0.262)	−0.186	(0.061)	0.203	(0.161)
GDP deflator (SPF, post ’92)	0.572	(0.246)	−0.381	(0.048)	−0.391	(0.094)
CPI inflation (SPF, post ’92)	0.272	(0.414)	−0.279	(0.175)	−0.555	(0.172)
Real GDP (SPF, post ’92)	0.601	(0.379)	−0.087	(0.135)	−0.584	(0.228)
GDP deflator (SPF, h=2)	0.266	(0.168)	−0.381	(0.044)	0.111	(0.054)
GDP deflator (SPF, finan.)	0.608	(0.261)	−0.377	(0.058)	−0.361	(0.083)
GDP deflator (SPF, nonfinan.)	0.295	(0.212)	−0.379	(0.039)	−0.293	(0.118)
HICP inflation (EASPF)	0.500	(0.436)	−0.169	(0.182)	**−0.669**	(0.665)
Real GDP (EASPF)	0.616	(0.226)	0.411	(0.170)	−0.905	(0.210)
CPI inflation (LIV)	1.184	(0.736)	−0.270	(0.077)	−0.193	(0.206)
Real GDP (LIV)	0.272	(0.201)	−0.325	(0.130)	−0.709	(0.371)

Estimates of b in equation (3), β in equation (2), and δ in equation (3), where the estimates of δ use the previous period's consensus outcome from the survey in question. LIV denotes the Livingston Survey, while EASPF refers to the Euro Area Survey of Professional Forecasters. All estimates are computed using year-on-year growth rates that have been derived using the latest available data at the time of the forecast. Shaded (dark or light) gray coefficients are significant at the five percent level. Robust (doubled-clustered) standard errors are used. Bold indicates a coefficient in which fewer than 50 time clusters have been estimated, and that is significant using the adjustment in Cameron et al. (2010). Samples: SPF(1970Q1–2020Q1), LIV(1993Q1–2020Q1), EASPF(2000Q1–2020Q1).

and Stock and Watson (2008), among others. To make our estimates in figure 2 comparable, all variables have been standardized, and have been signed such that an increase predicts higher inflation one year out.

On balance, we find that, although forecasters overreact to previous consensus forecasts from the SPF, the evidence for other public signals is more mixed. For example, the left-hand side panel in figure 2 shows that forecasters *overreact* with similar strength to the observation of respondents' consensus estimate from the *Livingston Survey* (Croushore, 1997). This is consistent with the *Livingston Survey* covering many of the same forecasters as the SPF. However, forecasters *underreact* to the information contained in the consensus outcome from the *Survey of Consumer Expectations* (Armantier et al., 2017), in addition to estimates of consumer expectations from the *Michigan Survey of Consumers* (Dominitz & Manski, 2003), although the latter is not statistically significant.[14]

The right-hand side panel in figure 2 confirms this picture of over- and underreactions in response to public signals other than measures of average expectations. When we estimate the relationship between individual inflation errors and nine common public signals of future inflation, we find significant overreactions to some (e.g., lagged outcomes, akin to extrapolation), but significant underreactions to others (e.g., changes to the exchange rate or the unemployment rate). A simple ANOVA exercise shows that the probability of all coefficients in figure 2 occurring by chance in the absence of over- or underreactions is less than 0.001.

Finally, two wider implications of our analysis are worth noting. First, the above analysis considers multiple public signals. However, our estimates do not attempt to directly estimate the relative weight placed on any specific signal compared to the rational expectations case. Such an exercise would require a full list of signals observed by forecasters, including those from private sources. Instead, our estimates explore the extent to which individual forecasts $f_{it}\pi_{t+h}$ are associated more or less with a public signal y_t than their rational counterpart $\mathbb{E}_{it}\pi_{t+h}$. Notwithstanding such concerns, online appendix table B.2 shows that a multivariate version of equation (3) still confirms the above picture of over- and underreactions.

Second, our findings of overreactions to consensus estimates are robust to concerns of *limited attention*. Although professional forecasters track developments in the above public signals closely, if they instead of the consensus signal y_t were to observe $z_{it} = y_t + u_{it}$, with $u_{it} \sim \mathcal{N}(0, \sigma_u^2)$, due to limited attention, then estimates of δ would be upward biased. This is for the same reason that noisy information leads to a positive b in regression equation (1). We, nevertheless, view large amounts of inattention to salient public signals, such as consensus, to be unlikely for the professional forecasters that comprise our sample.

D. *Alternative Estimates*

We obtain similar estimates to those in figures 1 and 2 beyond one-year-ahead inflation forecasts from the U.S. SPF. Table 1 and figures B.1 and B.2 summarize alternative estimates of equations (1), (2), and (3) using different variables and other expectational surveys.

First, to complement our benchmark results using GNP/GDP inflation forecasts, we consider forecasts of CPI inflation and real output growth (Real GDP) from the *Survey of Professional Forecasters* (table 1). The estimated coefficients for b in equation (1), β in equation (2), and δ in equation (3) using past consensus outcomes ($y_t = f_{t-1}\pi_{t+h}$) all have the same sign as our benchmark results, and are all statistically significant, with the exception of the CPI estimate

gap", since that would entail taking a stance on the structural determinants of deviations from the flex-price allocation.

[14]The Livingston Survey is a bi-annual survey that covers many different types of forecasters. It is the oldest continuous survey of forecaster's expectations. The Federal Reserve Bank of Philadelphia took responsibility for the survey in 1990. The Michigan Consumer Survey and the Survey of Consumer Expectations are monthly surveys of U.S. households.

of b and the output estimate for δ. Similar results to those in figure 1 also hold when we restrict the sample to after '92, when the Philadelphia Federal Reserve Bank took over the administration of the SPF. We also document similar patterns at a semi-annual forecast horizon ($h = 2$).

Second, we extend beyond the United States and consider professional forecasts for another geographic area, the Euro Area, as collected by the *ECB's Survey of Professional Forecasters* (Garcia, 2003). We once more find estimates of b, β, and δ using past consensus similar to those from the U.S. SPF. While the point estimate of β for output is positive, the uncertainty around this estimate is large because of the short sample that starts only in 2000. As we discuss below, our model in section IV can in any case also account for such observations.

Third, a large share of forecasters in the United States and Euro Area SPF comes from financial-sector institutions. Table 1 shows that our results carry over with equal force to the nonfinancial sector forecasters in the U.S. SPF (large private sector firms), as well as to the broader range of nonfinancial sector institutions that are part of the semiannual Livingston Survey. Table B.3 in the online appendix elaborates on these results.

Fourth, to further complement our baseline estimates, figures B.1 and B.2 summarize estimates of the under-/overreaction coefficient δ in equation (3), using alternative forecaster surveys and other public signals than previous consensus outcomes from the same survey. The estimates confirm our initial take-away from figure 2. For both the Euro Area SPF, the Livingston Survey, as well as U.S. SPF forecasts of CPI inflation, the overrevision of individual forecasts appears to be the product of both over- and underreactions to public signals ($\delta \lesseqgtr 0$). Hence, at a more general level, the documented overrevision of individual forecasts $\beta < 0$ is comprised of both over- and underreactions to public information ($\delta \lesseqgtr 0$).

Finally, tables B.4–B.6 in the online appendix contain further robustness checks. We document that the coincidence of over- and underreactions extend to cases where we consider median-individual estimates of equations (1), (2), (3), and that our results also extend to cases where we winsorize outlier observation (table B.5 and B.6). Further, table B.4 shows that if we drop forecaster i from the SPF consensus the overreaction to consensus documented above remains. This also holds when we exclude outlier observations. Lastly, consistent with the results in Clements (2018), table B.4 documents a negative correlation between individual errors and past consensus deviations. This will be important for later.

E. Summary and Discussion

In summary, our results suggest that *average* forecasts are consistent with models of noisy rational expectations with mean-squared-error preferences ($b > 0$). This confirms the results of Coibion and Gorodnichenko (2015). However, *individual* forecasts show patterns that strongly contradict such models. Specifically, forecasters systematically overreact, on average, to the news that they receive between subsequent survey rounds. This leads to too large forecast revisions relative to the noisy-rational expectations benchmark ($\beta < 0$). Consistent with this pattern of overall overrevisions, we find strong evidence of overreactions to a particular public signal that is salient in the context of professional forecasts, namely the consensus forecast from the previous round of the survey ($\delta < 0$). We, nevertheless, also find evidence of sizable underreactions to other public signals ($\delta > 0$). As we have argued in the introduction, and will show formally below, several prominent models of forecaster behavior, both rational and behavioral, struggle to explain this coincidence of over- and underreactions. The next section makes this explicit using a workhorse noisy information framework.

III. Rational and Behavioral Models

A variety of popular models of forecaster behavior are consistent with the under- and overrevision of forecasts at the average ($b > 0$) and individual level ($\beta < 0$), respectively. In this section, we show that several of the most prominent of such explanations are nevertheless inconsistent with the documented over- and underreaction to public information ($\delta \lesseqgtr 0$).

A. Model Environment

We outline a model that captures several popular environments used to describe economic forecasts. The model is comprised of a continuum of measure one of forecasters, indexed by $i \in [0, 1]$. Forecasters minimize the mean-squared error of their forecasts f_{it} of the random variable π_{t+h} drawn from an uniform distribution over the real line. At time $t = 1, 2, 3 \ldots$, all forecasters have the prior belief that $\pi_{t+h} \sim \mathcal{N}(\mu_{\mathrm{it}}, \tau_\pi^{-1})$ and observe two types of information.[15] Their own private information is summarized by the *private signal*

$$x_{\mathrm{it}} = \pi_{t+h} + \epsilon_{\mathrm{it}}^x, \quad \epsilon_{\mathrm{it}}^x \sim \mathcal{N}(0, \tau_x^{-1}), \tag{5}$$

where the noise terms ϵ_{it}^x are independent across time and of π_{t+h}, and $\mathbb{E}[\epsilon_{\mathrm{it}}^x \epsilon_{js}^x] = 0$ for all $j \neq i$ and $s \neq t$. The private signal of one forecaster is not observed by any other forecaster. In addition to their private information, all forecasters observe the *public signal*

$$y_t = \pi_{t+h} + \epsilon_t^y, \quad \epsilon_t^y \sim \mathcal{N}(0, \tau_y^{-1}), \tag{6}$$

where ϵ_t^y is independent across time, and of π_{t+h} and ϵ_{it}^x for all t and $i \in [0, 1]$. We note that this environment allows for rich heterogeneity in expectations, arising from both private

[15]We assume that all prior information is condensed into the signal $\mu_{\mathrm{it}} = \pi_{t+h} + \nu_{\mathrm{it}}$, where $\nu_{\mathrm{it}} \sim \mathcal{N}(0, \tau_\pi^{-1})$, observed in period $t - 1$. Hence, before the observation of period-t information, forecasters' beliefs about $\pi_{t+h} \sim \mathcal{N}(\mu_{\mathrm{it}}, \tau_\pi^{-1})$. We further have that $\mu_{\mathrm{it}} = f_{it-1}\pi_{t+h}$.

information and heterogenous prior forecasts ($\mu_{it} \neq \mu_{jt}$ for $j \neq i$). Finally, we assume that individual forecasts follow the *Generalized prediction rule*:[16]

$$f_{it}\pi_{t+h} = (1 - k_x - k_y)\mu_{it} + k_x x_{it} + k_y y_t, \tag{7}$$

where k_x and k_y may differ from the mean-squared optimal (rational) weights on private and public information, $k_x^\star \equiv \frac{\tau_x}{\tau_\pi+\tau_x+\tau_y}$ and $k_y^\star \equiv \frac{\tau_y}{\tau_\pi+\tau_x+\tau_y}$, respectively. The rational weight on the prior μ_{it} is $\frac{\tau_\pi}{\tau_\pi+\tau_x+\tau_y}$. As with the conditional expectation $\mathbb{E}[\pi_{t+h} \mid \mu_{it}, x_{it}, y_t]$, an individual forecaster in equation (7) updates her prior expectation $\mu_{it} = f_{it-1}\pi_{t+h}$ in response to private and public information, x_{it} and y_t, respectively. But, importantly, relative to the conditional expectation, the forecaster can both over- or underreact to private and public information.

Depending on the precise use of information, the forecasts from equation (7) predict specific values for the regression coefficients b in equation (1), β in equation (2), and δ in equation (3). We note that if $k_x = k_x^\star$ and $k_y = k_y^\star$, the coefficients b, β, and δ are all equal to zero. Proposition 1 summarizes two other important cases, which combined capture a popular set of alternative models.[17]

Proposition 1. *Let individual forecasts $f_{it}\pi_{t+h}$ follow the Generalized prediction rule equation (7).*

(i) Then, if $k_x \in (k_x^\star, 1)$ and $k_y = (1 - k_x)\frac{\tau_y}{\tau_\pi+\tau_y}$, so that

$$f_{it}\pi_{t+h} = (1 - k_x)\mathbb{E}[\pi_{t+h} \mid \mu_{it}, y_t] + k_x x_{it}, \quad k_x \in (k_x^\star, 1), \tag{8}$$

$b > 0$ in equation (1), $\beta < 0$ in equation (2), but $\delta = 0$ in equation (3).

(ii) Then, if $k_x = k\frac{\tau_x}{\tau_x+\tau_y}$ and $k_y = k\frac{\tau_y}{\tau_x+\tau_y}$ with $k \in (k^\star, 1)$, where $k^\star \equiv \frac{\tau_x+\tau_y}{\tau_\pi+\tau_x+\tau_y}$, so that

$$f_{it}\pi_{t+h} = \mu_{it} + k(\mathbb{E}[\pi_{t+h} \mid x_{it}, y_t] - \mu_{it}), \quad k \in (k^\star, 1), \tag{9}$$

$b > 0$ in equation (1), $\beta < 0$ in equation (2), but $\delta < 0$ in equation (3).

The first part of proposition 1 characterizes individual responses when forecasters over-emphasize private information $k_x \in (k_x^\star, 1)$. When forecasters attach more weight to private information than optimal, forecasters will, on average, overreact to the information that they receive between two periods. This leads to a negative correlation between individual errors, on the one hand, and individual revisions, on the other hand. Furthermore, this negative correlation coincides with an underrevision of the average forecast ($b > 0$). This is because forecasters with $k_x < 1$ still respond less to private information than the optimal reaction to the *average private signal* ($\int_0^1 x_{it} di = \pi_{t+h}$), which in this case equals one.[18]

However, while an increased weight on private information is consistent with our first two stylized facts, it leads to neither an over- nor an underreaction to public information. In fact, when $k_x > k_x^\star$ (or $k_x \neq k_x^\star$), errors remain uncorrelated with the public signal ($\delta = 0$).

The reason is that a regression of individual errors onto any public signal only considers whether that source of information is used to minimize forecast errors. It does not consider more broadly whether all sources of information, in general, are accurately employed. Although forecasters in equation (8) do not optimally use private information to minimize errors, conditional on this misuse, they still use public information efficiently. The expression $\mathbb{E}[\pi_{t+h} \mid \mu_{it}, y_t]$ enters in equation (8). This, in turn, leads to a δ coefficient that is equal to zero.[19]

The second part of proposition 1 considers a natural extension that simultaneously skews forecasters' use of private *and* public information away from their mean-squared optimal values. When k exceeds its optimal value ($k > k_\star$), forecasters in equation (9) overemphasize new information contained in private and public signals relative to their prior expectation. In this sense, forecasters with $k > k_\star$ overemphasize *all news* that is characteristic of updates relative to prior beliefs. When forecasters overreact to all information, the resulting forecasts from equation (9) can also be consistent with the documented behavior of forecast revisions ($b > 0$, $\beta < 0$). This occurs when $k \in (k_\star, 1)$. But, because forecasters overreact to all information such forecasts are also inconsistent with the documented underreaction to public signals ($\delta > 0$). Instead, such forecast always entail overreactions to public information ($\delta < 0$).

[16]Although we, for simplicity, adopt a framework in which inflation is drawn in each period, we note that our signal structure and prediction rule also carry over to other cases. For example, if inflation π_t follows an AR(1) process with persistence $\rho \in (0, 1)$, individual forecasts follow the recursion $f_{it}\pi_{t+k} = \rho^k f_{it}\pi_t$, where $k \geq 1$. In this case, equations (5) and (6) can still be assumed with $h = 0$. The rational expectation forecast is, furthermore, still a special case of equation (7) and rational errors remain uncorrelated with y_t.

[17]We note that we adjust for the bias caused by public information in our derivation of the regression coefficient b in Proposition 1. As mentioned in section II, because of the downward nature of this bias, our empirical findings of $b > 0$ are robust to the presence of public information (see also online appendix D).

[18]We note that an increased weight on private information $k_x \in (k_x^\star, 1)$ is also consistent with our results in online appendix table B.4, which documents a negative correlation between individual errors and past deviations of forecasts from consensus (Clements, 2018).

[19]Consider the forecast error that results from equation (8):

$$\pi_{t+h} - f_{it}\pi_{t+h} = \pi_{t+h} - k_x x_{it} - (1 - k_x)\mathbb{E}[\pi_{t+h} \mid \mu_{it}, y_t].$$

Taking conditional expectations based upon the public signal y then shows that

$$\begin{aligned}\mathbb{E}[\pi_{t+h} - f_{it}\pi_{t+h} \mid y_t] &= \delta \times y_t \\ &= (1 - k_x)(\mathbb{E}[\pi_{t+h} \mid y_t] - \mathbb{E}\{\mathbb{E}[\pi_{t+h} \mid \mu_{it}, y_t] \mid y_t\}) = 0,\end{aligned}$$

where the last equality follows from the *Law of iterated expectations*. Hence, despite the erroneous use of private information, errors remain uncorrelated with the public signal ($\delta = 0$).

B. Applications and Extensions

A variety of popular models of forecaster behavior fall within the cases described in proposition 1, where $b > 0$, $\beta < 0$, but $\delta \leq 0$. Below, we outline several of these.

Strategic diversification. Laster et al. (1999), Ottaviani and Sørensen (2006), and Marinovic et al. (2013) describe the market for professional forecasters as a winner-takes-all competition, where only the most accurate forecast is rewarded. As a consequence, in a symmetric equilibrium, all forecasters over-emphasize private information and follow equation (8) with $k_x > k_x^\star$.[20]

Reputational considerations. In Ehrbeck and Waldmann (1996), forecasters are rewarded based on their perceived accuracy. One set of forecasters has access to more precise private information than another. As a result, the set of forecasters that receive less precise information overreact to their private information in an attempt to mimic their more informed competitors, and follow equation (8) with $k_x > k_x^\star$. Their more informed competitors set $k_x = k_x^\star$. The average individual forecast thus follows equation (8) with $k_x > k_x^\star$ (online appendix C.1).[21]

Behavioral overconfidence. A considerable literature in psychology has documented that agents over-emphasize their own information (e.g., Moore & Healy, 2008). As discussed in, for example, Daniel et al. (1998), and more recently in Angeletos et al. (2021), such inherent overconfidence could provide a basis for overreactions to new information. Within our context, overconfident forecasters believe the precision of their private information to be higher than it actually is. Their forecasts thus follow equation (8) with $k_x \in (k_x^\star, 1)$. We return to how a suitably adjusted notion of behavioral overconfidence can capture our stylized facts in section IV.

Models of generalized overreactions. A candidate explanation for the overreaction to individual information ($\beta < 0$) and consensus expectations ($\delta < 0$) that we have documented are models of *generalized overreactions*. This includes Bordalo et al.'s (2018a) theory of *diagnostic expectations* and Evans and Honkapohja (2012)'s theory of *excess Kalman Gain learning*. In the former case, forecasters overreact to all new information, because it is perceived to be diagnostic (or representative) of updates relative to prior information. In the latter case, forecasters instead overreact to increase their speed of learning. Within our framework, these models are captured by equation (9) with $k \in (k^\star, 1)$ (online appendix C.2).

Underreactions and rational inattention. We close this list by noting that several other, prominent models of forecaster behavior fall within the cases described in proposition 1, but where $k_x \in (0, k_x^\star)$ or $k \in (0, k^\star)$.[22] As a result, these models cannot explain the documented overrevision of individual forecasts ($\beta < 0$). Finally, we note that models of rational inattention (e.g., Sims, 2003), or other rational models of limited attention, are likewise inconsistent with $\beta < 0$. This is because forecasts from these models equal conditional expectations, and hence satisfy the Law of Iterated Expectations.[23]

The above examples have shown that several prominent models of forecaster behavior are consistent with under- and overrevisions of expectations at the average ($b > 0$) and individual level ($\beta < 0$), respectively. However, none of these models have been simultaneously consistent with the documented over- and underreaction to public information ($\delta \lesseqgtr 0$). This insight extends beyond the specific applications considered above.

Online appendix C.3 analyzes a more general model, where strategic incentives skew the optimal use of information away from its mean-squared optimal value. This appendix shows that, despite flexible strategic interactions, errors remain uncorrelated with public information. This result extends to cases with a common noise component in private information. Online appendix C.4 shows that our results also extend to circumstances where trembling-hand noise drives a wedge between *reported* estimates and *actual* expectations.

Clearly, extensions or combinations of the above environments could potentially alter the prediction listed in proposition 1. But, at this point, it is worth summarizing why these models fail to match the data. At its heart, the reason is that to explain the survey data forecasters have to flexibly misperceive public information. As Part (ii) of proposition 1 shows, forecasters cannot, for example, always place an excessive weight on public information. Whatever misperception we

[20]To see why, consider an individual forecaster who sets $k_x = k_x^\star$. Increasing the weight on private information ($k_x > k_x^\star$) leaves the probability of winning the contest approximately unchanged (as the posterior is flat at the conditional expectation). But more weight on private information also (in expectation) strictly reduces the mass of other forecasters that makes the same forecast. In equilibrium, all forecasters therefore choose to follow equation (8) and set k_x such that $k_x \in (k_x^\star, 1)$ (see, e.g., proposition 4 in Ottaviani & Sørensen, 2006 and proposition 1 and Corollary 1 in Marinovic et al., 2013).

[21]See the results on p. 24 of Ehrbeck and Waldmann (1996). Online appendix C.1 extends their model to explicitly account for public information. We assume that forecasters as well as clients observe the public signal y_t in equation (6). We summarize all initial information in the individual-specific prior μ_{it}. With the exception of these modifications all details are as in Ehrbeck and Waldmann (1996).

[22]For example, Graham (1999), Lamont (2002), and Ottaviani and Sørensen (2006) describe models in which forecasters all have a rational incentive to herd, as in Scharfstein and Stein (1990). Hirshleifer et al. (2011) instead detail a model in which security analysts for behavioral reasons underreact to information. All of these explanations feature either $k_x \in (0, k_x^\star)$ in equation (8) or $k \in (0, k^\star)$ in equation (9).

[23]Let $x_{\text{it}}^\star$ denote the optimal signal observed by a capacity-constrained agent with entropy attention cost. Following Maćkowiak and Wiederholt (2009), $x_{\text{it}}^\star$ follows equation (5) but with a precision $\tau_x^\star \neq \tau_x$ that depends upon the capacity constraint. The agent's forecast equals $f_{\text{it}}\pi_{t+h} = \mathbb{E}[\pi_{t+h} \mid \mu_{\text{it}}, x_{\text{it}}^\star]$. But then the exact same steps as those taken in the proof of Proposition 1 show that $\beta = 0$, because of the Law of iterated expectations.

consider has to result in both too much as well as too little weight on public signals. The next section shows that a natural candidate for such flexible misperceptions arises from forecasters' potentially incorrect views about other's information.

IV. Absolute and Relative Overconfidence

In this section, we show that a simple model in which forecasters are overconfident in the precision of their own information (both relative to the truth and relative to their perception of others) can account for all three stylized facts. The next section then explores the potential of our model to also quantitatively match the magnitude of our empirical estimates.

A. *Overconfidence and Public Information*

We build our model of expectations from first principle starting with the well-documented overconfidence heuristic. In their overview of behavioral finance, De Bondt and Thaler (1985) state that "perhaps the most robust finding in the psychology of judgement is that people are overconfident" (p. 6). In particular, we call overconfident those individuals that are not only overconfident in the precision of their own information but also wrongly think that their information is better than others. We therefore merge the two related but distinct notions of overconfidence commonly used in the psychology literature (Moore & Healy, 2008). We refer to the first type as *absolute overconfidence* and the second type as *relative overconfidence* (Benoît et al., 2015). Notice that it is the second, relative aspect of overconfidence that differentiates the notion of overconfidence studied here from that explored in section III. To motivate these assumptions, we briefly return to the survey data.

Panel a in table 2 uses data on individual-level density forecasts of one-year-ahead inflation from the U.S. SPF. It shows that respondents' stated accuracy of their one-year-ahead inflation forecasts exceeds their actual accuracy by a sizable amount. The estimated *coverage ratio* of respondents' 95% confidence interval, which describes the percentage of times when inflation outcomes fall inside an individual respondent's confidence interval, is only between 72% and 84%, depending on the estimation method. Closely related, Griffin and Tversky (1992) show that such absolute overconfidence tends to be more prevalent for forecasters that are faced with prediction tasks that are characterized by a large judgment component and delayed feedback, such as economic forecasters.[24]

Panel b in table 2 uses the recent survey on firm managers' higher-order expectations of one-year-ahead inflation undertaken by Coibion et al. (2021) in New Zealand. Consistent with absolute overconfidence in individual-specific information, Coibion et al. (2021) document that the cross-sectional standard deviation of respondents' inflation forecasts is too large when compared to the predictions from simple noisy rational expectation models. However, crucially, Panel b also shows that the elicited higher-order moments from the survey combine with their first-order counterpart to imply substantial relative overconfidence (Appendix A.6). In particular, we can use the information about individual views of other's information in the data on uncertainty and disagreement about consensus expectations. This suggests that respondent i's perception of respondent $j \neq i$'s uncertainty of future inflation is, on average, much above that of her own, consistent with the hallmark of relative overconfidence. Indeed, respondents' estimates suggest that individuals believe, on average, that their expectations are 45% more accurate than their competitors.

Finally, our model accounts for the fact that most public signals are *endogenous*. A central feature of the information landscape that people observe is that most of it reflects the realized choices of others. This is true whether one considers data releases on past inflation or output, the observation of asset or goods prices, or the observation of previous period's consensus estimate. Because of this endogeneity of public signals, any *equilibrium* model of expectation

TABLE 2.—OVERCONFIDENCE IN SURVEY DATA ON EXPECTATIONS

Panel a: Coverage ratio of forecasts

	Confidence interval	
Estimation method	95%	66%
Density implied	0.85**	0.59*
Giordani and Söderlind (2003)	0.72**	0.44**

The table uses SPF density forecasts for one-year-ahead GDP deflator inflation. The table shows the coverage ratio (the fraction of cases when actual inflation is inside a forecaster's confidence band). If forecasters are rational a 95% confidence band will contain the true but unknown value 95% of the times. The confidence bands are derived assuming a normal distribution and are calculated as: mean of individual inflation densities ± critical value × standard deviation. Actual inflation is measured as the percentage change in the index (annual-average) in Q4 of each year. The significance of differences between the nominal confidence level and the actual are assessed using Christoffersen's (1998) test. $^{*}p < 0.1$, $^{**}p < 0.05$, and $^{***}p < 0.01$. The sample period is 1981Q1 to 2018Q4. For reference, the table also includes estimates from Giordani and Söderlind (2003).

Panel b: Uncertainty and dispersion of higher-order expectations

	Uncertainty	Std. dev.	Implied Higher-order Unc.	
	(1)	(2)	(1)	(2)
Inflation one-year-ahead	1.11	3.06	1.76	1.80
Consensus inflation expectation	0.89	2.43	–	–

The first column shows data on the average self-reported uncertainty of one-year ahead inflation expectations, as well as the average self-reported uncertainty about the consensus (average) expectation from the same survey. The second column shows the cross-sectional standard deviation (disagreement) of point forecasts of inflation and consensus. The data in the first two columns are taken from table II [initial wave] in Coibion et al. (2021). The last two columns instead use the data in the first two columns to compute respondents' perception of other respondents' uncertainty of future inflation ("implied higher-order uncertainty"; appendix A.6), measured as 1/(precision of prior + perceived precision of others' information). The third column uses the data in column one to compute this value, while the fourth column uses the data on disagreement in column two. Uncertainty is in all cases measured in terms of standard deviation. We throughout assume that the unconditional variance of inflation, the inverse of the prior precision, equals that realized post-1985 (Great Moderation) in New Zealand, the country to which the Coibion et al. (2021) survey pertains.

[24]Other prominent examples of overconfidence include the stated precision of forecasts produced by financial market traders, the certainty in the diagnosis of severe illnesses by physicians, and the probability of a positive verdict by procedural lawyers. See, for example, Einhorn (1980) and the summaries in, for example, Moore and Healy (2008).

formation requires an assumption about individuals' views about the precision of others' information. Rational expectations commonly solve this issue by imposing the symmetry assumption that others' information is equal in quality to one's own. Relative overconfidence, by contrast, imposes the empirically motivated "better than others" perception.

B. *Environment with Overconfidence*

We modify our previous environment from section III. We assume that inflation π_{t+h} is drawn from the normal distribution $\pi_{t+h} \sim \mathcal{N}(0, \underline{\tau}_\pi^{-1})$. At the start of period $t-1$ and t, each forecaster $i \in [0, 1]$ receives the private signal $x_{i\tau}$ about the fundamental π_{t+h},

$$x_{i\tau} = \pi_{t+h} + \epsilon_{i\tau}^x, \quad \epsilon_{i\tau}^x \sim \mathcal{N}(0, \tau_x^{-1}), \tag{10}$$

where $\tau = \{t-1, t\}$ and $\epsilon_{i\tau}^x$ is independent of π_{t+h} with $\mathbb{E}[\epsilon_{i\tau}^x \epsilon_{js}^x] = 0$ for all $j \neq i$ and s.[25] We introduce the period $t-1$ signal to later allow the public signal that forecasters observe to depend on the previous period's consensus expectation. All forecasters exhibit absolute and relative overconfidence. They believe the precision of their private signals equals $\tau_x' > \tau_x$, and thus to be greater than the truth (absolute overconfidence). At the same time, forecasters also believe that other forecasters' private signals have a precision smaller than their own (relative overconfidence) equal to $\hat{\tau}_x < \tau_x'$. We make no assumptions about the relative size of $\hat{\tau}_x$ and τ_x. We note that the observation of x_{it-1} at time $t-1$ results in a period-t prior of π_{t+h} of the exact form used in section III (see further below).

At the start of period t, each forecaster, in addition, observes the *endogenous* public signal

$$y_t = \alpha_1 \pi_{t+h} + \alpha_2 f_t \pi_{t+h} + \epsilon_t^y, \quad \epsilon_t^y \sim \mathcal{N}(0, \tau_y^{-1}), \tag{11}$$

where $\alpha_j \geq 0$ for $j = \{1, 2\}$, and ϵ_t^y is independent of π_{t+h} and ϵ_{is}^x for all $i \in [0, 1]$ and s.[26] The key difference between the public signal in equation (11) and that explored in equation (6) is the endogeneity of the signal to average individual expectations $f_t \pi_{t+h}$. For example, when $\alpha_1 = 0$ and $\alpha_2 = 1$, equation (11) directly becomes the consensus (average) forecast of inflation from the previous period. Veldkamp (2011) summarizes the importance of public signals of the form equation (11) for the social value of public information, the benefits of social learning, and the volatility of asset prices and business cycles, among others.

We proceed in two steps. We first derive individual expectations of inflation π_{t+h} in period $t-1$ and t, and show how relative overconfidence causes forecasters to flexibly misperceive the public signal y_t. We then provide a set of sufficient conditions for individual expectations to be consistent with all three of our stylized facts.

Consider forecaster i's expectation of π_{t+h} in period $t-1$:

$$f_{it-1} \pi_{t+h} = \upsilon x_{it-1}, \tag{12}$$

where $\upsilon \equiv \frac{\tau_x'}{\tau_x' + \underline{\tau}_\pi}$ exceeds the mean-squared optimal weight on the private signal $\upsilon^\star \equiv \frac{\tau_x}{\tau_x + \underline{\tau}_\pi}$, because of forecaster i's (absolute) overconfidence in her private information. Importantly, the coefficient υ also exceeds the weight that the forecaster believes others place on their private information (because of relative overconfidence), equal to $\hat{\upsilon} \equiv \frac{\hat{\tau}_x}{\underline{\tau}_\pi + \hat{\tau}_x}$.

Let $\mu_{\text{it}} \equiv f_{it-1} \pi_{t+h}$ denote forecaster i's prior expectation at the start of period t with perceived precision $\tau_\pi \equiv \underline{\tau}_\pi + \tau_x'$. To derive forecaster i's period-t expectation, we first need to differentiate between two different public signals: (i) the *realized public signal* y_t, and (ii) the *perceived public signal* $\hat{y}_t$. The former measures the actual signal in equation (11),

$$y_t = \alpha_1 \pi_{t+h} + \alpha_2 \int_0^1 f_{it-1} \pi_{t+h} di + \epsilon_t^y = \eta \pi_{t+h} + \epsilon_t^y, \tag{13}$$

where $\eta \equiv (\alpha_0 + \alpha_1 \upsilon) > 0$. The latter, by contrast, measures the public signal that forecasters believe they observe when confronted with observations of y_t,

$$\hat{y}_t = \hat{\eta} \pi_{t+h} + \epsilon_t^y, \tag{14}$$

where $\hat{\eta} \equiv (\alpha_0 + \alpha_1 \hat{\upsilon}) > 0$. Notice that the signals y_t and $\hat{y}_t$ differ only because of forecasters' misperception about the overconfidence of others ($\eta > \hat{\eta}$); that is, because all forecasters attach a weight of $\upsilon > \hat{\upsilon}$ to private information in equation (12). This shows how relative overconfidence boils down to a simple one-parameter deviation from rational expectations.

We are now ready to state forecasters' period-t expectation. Combining the public signal in equation (13) with forecasters' perception about it in equation (14), as well as with the period-t private signal in equation (10), shows that

$$f_{\text{it}} \pi_{t+h} = (1 - k_x - k_y) \mu_{\text{it}} + k_x x_{\text{it}} + k_y \times \frac{1}{\hat{\eta}} y_t \tag{15}$$

$$= (1 - k_x) \mathbb{F}[\pi_{t+h} \mid \mu_{\text{it}}, y_t] + k_x x_{\text{it}}, \tag{16}$$

where $k_x \in (k_x^\star, 1)$, $k_y \lesseqgtr k_y^\star$, and $k_x^\star$ and $k_y^\star$ once more denote the mean-squared optimal weight on private and public information, respectively (appendix A.2).

We conclude from equation (15) that forecasters' expectations are a special case of those from the generalized prediction rule in equation (7). Equation (16) shows that these expectations can also be recast in a form similar to that studied in case (i) in proposition 1. The difference being that the conditional expectation $\mathbb{E}[\pi_{t+h} \mid \mu_{\text{it}}, y]$ in equation (16)

[25]Hence, at time t, a forecaster receives two private signals: one for π_{t+h} and one for π_{t+h+1}. To avoid complicating the notation further, we do not add an additional subscript on $x_{i\tau}$ to keep a track of the distinct private signals observed at time t. We can do so because of the independence of π_{t+h} and π_{t+h+1}. Neither of our results, however, depend on this independence feature; it merely simplifies the exposition.

[26]We restrict the sign of α_1 and α_2 to avoid having to always separate between positive and negative signals of the fundamental in our discussions. Neither of our main results depend importantly on this assumption.

is replaced with the overconfident forecast $\mathbb{F}[\pi_{t+h} \mid \mu_{\text{it}}, y]$ that accounts for the misperception of the public signal; that is, the conditional expectation of π_{t+h} based on μ_{it} and y_t, but where a forecaster perceives y_t to be governed by equation (14) instead of equation (13).

C. Over- and Underreactions to Public Information

Because of the misperception of the public signal, a correlation naturally arises between individual errors, on the one hand, and the public signal, on the other hand. Taking conditional expectations of forecaster i's error based upon the *realized* public signal y_t shows that

$$\begin{aligned}\delta \times y &= \mathbb{E}[\pi_{t+h} - f_{\text{it}}\pi_{t+h} \mid y_t] \\ &= (1-k_x)(\mathbb{E}[\pi_{t+h} \mid y_t] - \mathbb{E}[\mathbb{F}[\pi_{t+h} \mid \mu_{\text{it}}, y_t] \mid y_t]) \\ &= (1-k_x)\mathbb{E}\{\mathbb{E}[\pi_{t+h} \mid \mu_{\text{it}}, y_t] \\ &\quad - \mathbb{F}[\pi_{t+h} \mid \mu_{\text{it}}, y_t] \mid y_t\} \neq 0, \qquad (17)\end{aligned}$$

where we have used the expectation in equation (16) and the Law of Iterated Expectation to arrive at the second and third condition, respectively. Unlike with case (i) in proposition 1, the Law of iterated expectations in equation (17) does not imply orthogonality between individual errors and public information. This is because $\mathbb{E}[\pi_{t+h} \mid y_t] \neq \mathbb{E}\{\mathbb{F}[\pi_{t+h} \mid \mu_{\text{it}}, y_t] \mid y_t\}$. The misperception of the public signal breaks the implication of the Law of Iterated Expectations that forecast errors are orthogonal to public information. Proposition 2 computes an expression for the over- and underreaction coefficient δ in equation (17).

Proposition 2. *The over- and undereaction coefficient* δ *in equation (3) equals*

$$\delta = \Delta(\kappa^\star - \hat{\kappa}), \qquad (18)$$

where $\Delta \in \mathbb{R}_+$, $\kappa^\star \equiv \frac{\eta^2\tau_y}{\underline{\tau}_\pi + \eta^2\tau_y} \times \frac{1}{\eta}$ *denotes the rational weight on the public signal* y_t *in* $\mathbb{E}[\pi_{t+h} \mid y_t]$*, while* $\hat{\kappa} \equiv \frac{\hat{\eta}^2\tau_y}{\underline{\tau}_\pi + \hat{\eta}^2\tau_y} \times \frac{1}{\hat{\eta}}$ *denotes the corresponding misperceived weight.*

Intuitively, how forecasters respond to a public signal, such as past consensus outcomes, depends on their views about its precision (conditional variance) and its interpretation (conditional mean). Relative overconfidence causes forecasters to mistake both. On the one hand, it causes forecasters to underestimate the precision of public signals. The realized public signal y_t in equation (13) is more precise than the perceived public signal $\hat{y}_t$ in equation (14). The precision of the former is $\eta^2\tau_y$, while the precision of the latter is only $\hat{\eta}^2\tau_y$, where $\eta^2\tau_y > \hat{\eta}^2\tau_y$ since $\upsilon > \hat{\upsilon}$. The dismissal of other forecasters' information straightforwardly leads forecasters to *underreact* to the public signal ($\delta > 0$ as it causes $\kappa^\star > \hat{\kappa}$). On the other hand, relative overconfidence also causes forecasters to over-infer movements in fundamentals from public signals. The realized public signal y_t loads onto the fundamental π_{t+h} with η in equation (13), while the perceived public signal $\hat{y}_t$ only loads onto the fundamental with $\hat{\eta} < \eta$ in equation (14). Hence, a movement of $d\pi_{t+h} > 0$ in the fundamental causes forecasters to, all else equal, believe in a movement equal to $(\eta/\hat{\eta})d\pi_{t+h} > d\pi_{t+h}$, based on the observation of the public signal alone. This misinterpreation of the public signal, in turn, leads forecasters to *overreact* to its realizations. When forecasters overinfer values of the fundamental from observations of the public signal, they all else equal attach more weight to it than warranted ($\delta < 0$ as it causes $\kappa^\star < \hat{\kappa}$).

Depending on the relative strength of these effects, proposition 2 shows that both under- and overreactions to a public signal can arise from individuals' dismissal of other's private information. Indeed, equation (18) provides the condition for $\delta \lesseqgtr 0$.

We close this subsection with two additional observations that follow from proposition 2 (online appendix A.4). First, we note that *overreactions (underreactions)* to public signals naturally arise when the public signal that forecasters observe is sufficiently precise (imprecise). Equation (18) shows that $\lim_{\tau_y \to 0} \delta > 0$ while $\lim_{\tau_y \to \infty} \delta < 0$. In section V, we relate this finding to our empirical estimates of δ for different public signals in figure 2. Second, we note that when forecasters believe others' information is poor $\hat{\tau}_x \to 0$, equation (18) shows that underreactions always occur ($\delta > 0$). This can provide a lens through which to interpret some of our estimates using alternative consensus estimates in section II.

D. Data-consistent Expectations

Unlike the models in section III, the expectations in equation (16) can be consistent with all three stylized facts documented in section II. We show this concretely by focusing on our results in figure 1 ($b > 0$, $\beta < 0$, and $\delta < 0$), where we consider previous period's consensus estimate as the relevant public signal ($\alpha_1 = 0$ and $\alpha_2 = 1$). Section V explores the quantitative potential of our model to also match the magnitude of the empirical estimates.

Proposition 3. *Suppose* $\alpha_1 = 0$ *and* $\alpha_2 = 1$*, such that the public signal* y_t *corresponds to previous period's consensus estimate, and consider individual* $i \in [0, 1]$*'s forecast*

$$f_{\text{it}}\pi_{t+h} = (1-k_x)\mathbb{F}[\pi_{t+h} \mid \mu_{\text{it}}, y_t] + k_x x_{\text{it}}, \quad k_x \in (k_x^\star, 1). \qquad (19)$$

If $\eta^2\tau_y > \max(\chi, 1)\underline{\tau}_\pi$*, where* $\chi \equiv \frac{\tau_\pi \hat{\eta}}{\tau_x'(1-\hat{\eta})}$*, then there exists* $c_0, c_1 \in \mathbb{R}_+$ *such that, for* $\epsilon > 0$ *and* $\tau_x' = \tau_x + c_0\epsilon$ *and* $\hat{\eta} = \eta - c_1\epsilon$*, the coefficients satisfy* $\beta < 0$*,* $\delta < 0$*, and* $b > 0$*.*

Proposition 3 combines the insights of propositions 1 and 2. The first and second result in proposition 3 ($b > 0$ and $\beta < 0$) resemble those in case (i) of proposition 1. On the one hand, because of the dispersion in private signals, the average information across forecasters is more precise than any individual's. This, in turn, causes average forecasts to underreact to the average information observed ($b > 0$). On the other hand, despite these underreactions at the average level, at the individual level, forecasters overrevise their expectations ($\beta < 0$). This is once more in part due to forecasters' overconfidence in their own private information.

However, where proposition 3 differs from case (i) of proposition 1 is that forecasters also overreact to the past consensus outcomes ($\delta < 0$). These overreactions occur because forecasters' perceived and actual weight on private information are sufficient to ensure that the *perceived underresponsiveness of consensus* dominates its *perceived under-precision*. The condition $\eta^2\tau_y > \max(\chi, 1)\underline{\tau}_\pi$ ensures that relative overconfidence $\hat{\eta} < \eta$ delivers $\hat{\kappa} > \kappa^\star$ in equation (18) in proposition 2. Combined with the dispersion and overconfidence in private information, this then ensures that the expectations from equation (16) are consistent with all three stylized facts documented in panel b in figure 1.

V. Quantitative Implications

We have shown how our model of overconfidence can be qualitatively consistent with stylized facts about individual forecasts. Although our model is simple, in this subsection, we explore the capacity of the model to also quantitatively match the survey data. We also test several key implications of our model, and discuss its economic consequences.

A. *Model Calibration*

We use a simulated method of moments procedure to choose parameter values. Normalizing the variance of inflation to one and employing the restriction that $\hat{\tau}_x = \tau_x$, identification of the three parameters τ_x, τ_x', and τ_y requires at least three target moments. We choose the individual overrevision and overreaction coefficients β in equation (2) and δ in equation (3), respectively, documented in figure 1. We choose the previous consensus expectation as the benchmark public signal because its structure is simple and known [$\alpha_1 = 0, \alpha_2 = 1$ in equation (11)], and because its only relationship with future inflation is that of aggregating others' information. We then later show that the calibrated model also matches dimensions of the responses of individual errors to other public signals than consensus. Finally, we also include the estimate of information frictions b to our list of target moments. In particular, to account for the special feature that our baseline model only has one public signal, and not numerous as used by professional forecasters, we target the bias-adjusted measure of information frictions from Goldstein (2021). This estimate, in effect, bias adjusts the b coefficient in equation (1) for the presence of public information, and hence provides a more comparable estimate of the extent of information frictions to that of our model.[27] The criterion we choose to minimize is the sum of absolute deviations of target moments from model simulated moments.

Online appendix E presents the results for one-year ahead inflation, where for ease of interpretation we report the square root of the precision, the inverse of the standard deviation. Our model is able to capture all three data moments well. We estimate private signals to be rather noisy ($\sqrt{\tau_x} = 0.41$) and the noise in consensus to be small ($\sqrt{\underline{\tau}_\pi} = 4.74$). At a level of overconfidence that increases the square-root of the perceived precision of private signals by somewhat more than two ($\sqrt{\tau_x'} = 0.95$), the model predicts accurately the overrevision of individual forecasts β and the overreaction to past consensus realizations δ ($\beta^{\text{model}} = -0.18$ and $\delta^{\text{model}} = -0.19$). This is consistent with our previous discussion, which showed that the combination of a precise consensus and meaningful overconfidence, all else equal, makes overreactions more pervasive. The model also matches the level of information frictions well, although it entails somewhat too high information frictions ($b^{\star,\text{model}} = 0.58$ vs. $b^{\star,\text{data}} = 0.41$). Finally, online appendix E shows that our estimates also capture well the nontargeted Clements (2018) regression of individual errors onto consensus deviations.

B. *Model Evaluation*

Estimates of overconfidence. The calibrated model entails a noticeable degree of overconfidence. We next turn to how the implied estimates of absolute and relative overconfidence match those from survey data.

Estimates of absolute overconfidence. The individual density forecasts of one-year-ahead inflation, available in the U.S. SPF, allow us to evaluate whether the implied degree of absolute overconfidence from our model is reasonable. Panel a in table 3 presents this comparison in the form of *coverage ratios,* describing the percentage of times when actual inflation outcomes fall inside an individual forecaster's 95 (or 66) percent confidence band. Panel a contrasts the coverage ratios implied by our benchmark parameters with those that are estimated from U.S. SPF data in table 2. On balance, the implied degree of absolute overconfidence captures well that in the U.S. SPF data. Forecasters' 95 percent confidence band has a coverage ratio of only 70%, consistent with a sizable amount of absolute overconfidence. This matches the magnitude of the U.S. SPF estimate. In fact, Giordani and Söderlind (2003) find similar degrees of absolute overconfidence to those implied by our model estimates, using a somewhat more advanced estimation method to deduce individual confidence bands from reported forecast densities. We view the estimates in panel a in table 3 as important auxiliary evidence that corroborates our first main assumption of absolute overconfidence.

[27] Online appendix D provides details on the Goldstein (2021)-adjustment of the Coibion and Gorodnichenko (2015) estimate of information frictions in equation (1). In particular, the Goldstein (2021)-adjustment produces an unbiased estimate that is positively proportional to the bias-adjusted b coefficient in equation (1).

TABLE 3.—OVERCONFIDENCE IN SURVEY AND MODEL DATA

Panel a: Coverage ratio of forecasts

	Confidence level	
Confidence bands	95%	66%
SPF density implied	0.85**	0.59*
Giordani and Söderlind (2003)	0.72**	0.44**
Model implied	0.70	0.41

The table shows the implied coverage ratio. The confidence bands from the SPF are derived assuming a normal distribution and are calculated as: mean of individual density forecast ± critical value × standard deviation. Actual inflation is measured as the percentage change in the GDP Deflator (annual average). The significance of differences between the nominal confidence level and the actual are assessed using Christoffersen's (1998) test. $^{*}p < 0.1$, $^{**}p < 0.05$, and $^{***}p < 0.01$. The sample is from 1981Q1 to 2018Q4.

Panel b: Ratio of higher-order uncertainty to first-order uncertainty

	Higher-order uncertainty	
	(1)	(2)
Survey Data	1.59	1.62
Model Implied	1.28	1.28
Share (model of data)	0.81	0.79

The table estimates the ratio of higher-order uncertainty to first-order uncertainty in the survey data from Coibion et al. (2021), using the implied estimates in table 2. The table then compares these estimates to the model-implied counterparts. The model-implied results use the estimates of τ_x' and $\hat{\tau}_x = \tau_x$ from online appendix E. An absence of relative overconfidence results in a ratio of one. Uncertainty is measured in units of standard deviation. We express the results in terms of share of the data in the final row of the table to account for the different volatilities of inflation in New Zealand and the United States, and to account for our model only being calibrated to U.S. data. Consistent with table 2, columns denoted with a 1 use survey data on uncertainty about consensus to estimate higher-order uncertainty. Columns denoted with a 2 instead use survey data on the cross-sectional dispersion of forecasts of consensus to estimate higher-order uncertainty.

Estimates of relative overconfidence. Reported forecast densities can be used to assess the extent of *absolute overconfidence*, or the perceived precision of forecasters' information relative to the truth. In contrast, to assess the extent of *relative overconfidence* requires information about forecasters' perception of other forecasters' uncertainty (or expectations). This is typically not available in expectational survey data. The only exceptions that we are aware of are the survey of New Zealand firm managers, conducted by Coibion et al. (2021) and discussed in table 2, and the German ZEW survey (Köhler & Schmidt, 2021). The former asks its participants about their uncertainty about the consensus estimate of future inflation; the latter, by contrast, asks its respondents every month for their best forecast of the consensus estimate of an aggregate index of German economic activity.

Panel b in table 3 shows the model-implied estimates of the ratio of higher-order uncertainty (i.e., a forecaster's estimate of another forecaster's uncertainty about future inflation) to first-order uncertainty of future inflation, using our baseline parameters. An absence of relative overconfidence results in a ratio of one. The table compares these estimates to those implied by the Coibion et al. (2021) survey reported in table 2. Consistent with relative overconfidence, our model estimates show that forecasters perceive their own expectations to be around 30 percent more accurate than their competitors. As a result, our model accounts for around 80% of the relative overconfidence implied by the Coibion et al. (2021) data. That said, clearly, the implied estimates from our model, based on U.S. SPF data, are not fully comparable to those from the Coibion et al. (2021) survey, because of differences in respondent types (professional forecasters vs. managers) and countries covered (United States vs. New Zealand). Notwithstanding these discrepancies, the fact that the implied estimates in Panel b are of a similar magnitude is comforting, and the table does provide independent validation of our second main assumption of relative overconfidence.

Finally, online appendix G uses the time series data on higher-order expectations of economic activity, available from the ZEW survey, to directly estimate the actual and perceived weight on private information. Consistent with relative overconfidence, we estimate the actual weight on private information υ to be around twice the perceived weight attached by others $\hat{\upsilon}$, although the difference is only borderline statistically significant when accounting for outlier observations. In the calibrated model, the weight on private information υ is around three times the perceived weight.[28] We view these estimates, although pertaining to another variable and country, as lending further support to our assumption of relative overconfidence.

Heterogeneity in responses to public information. We revisit the evidence in figure 2, documenting heterogeneous responses to public information, ranging from over- to underreaction ($\delta \lesseqgtr 0$). In particular, we analyze how the over- and underreaction coefficient δ from our model changes with respect to the precision of public information. We then compare these predictions to estimates in the survey data.

The left-hand panel in figure 3 computes the model-implied estimates of δ as a function of the precision of the consensus signal in equation (13). The figure highlights two features of our model. First, all else equal, forecasters overreact more strongly to more precise signals (see also appendix A.4). Second, the precise parameters of the model determine the magnitude of over- and underreactions for any given precision of public information. The right-hand panel of figure 3, by contrast, returns to the survey data. It illustrates the relationship between the precision of different public signals of one-year ahead inflation and the over- and underreaction coefficient δ in equation (3), using our estimates in figure 2 but for a common sample. In line with the prediction of our model, we observe stronger overreactions to more precise signals. We note that the range of values in the left- and right-hand panels do not necessarily overlap, as the right-hand panel uses estimates employing other public signals than consensus. Overall, the results in figure 3 lend credence to the notion that an important determinant of over- and underreactions to public signals is the noisiness of the signal in question.

[28]We note that the estimated weight on private information in the ZEW is smaller than the model-implied estimate. The estimate from the ZEW data is roughly in line with that backed out by Coibion et al. (2021). This is consistent with the presence of several, additional public signals, beyond consensus estimates.

FIGURE 3.—OVERREACTION AND THE PRECISION OF PUBLIC SIGNALS

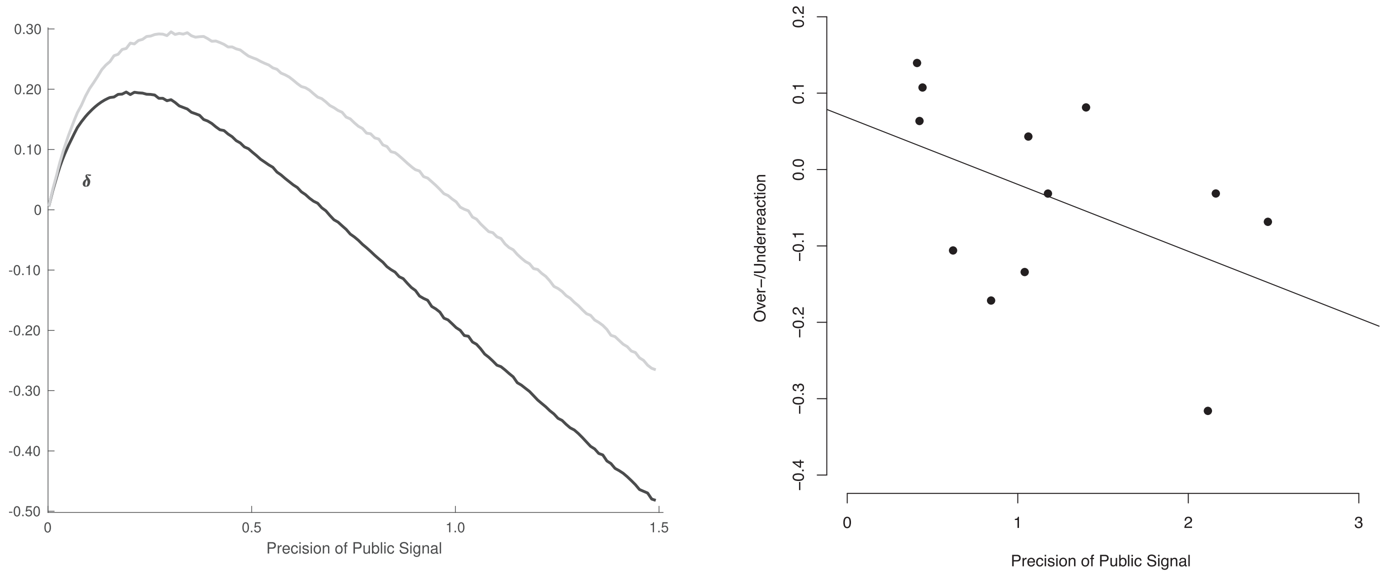

The left panel illustrates model-implied variations in δ as a function of the precision of consensus $\underline{\tau}_\pi$ relative to its calibrated value (black line). A value of one on the horizontal axis, therefore, corresponds to a precision of public information equal to that in online appendix E. The gray line decreases the precision of private information τ_x and forecasters' beliefs about it τ_x' by 25%. The right hand panel shows the estimates of δ for different public signals (along the vertical axis), using the variables from panels a and b in figure 2, as a function of the signals' estimated precision (along the horizontal axis). Consistent with section IV, we estimate the precision of public signals as the inverse of the variance of an error term. For consensus signals of the same forecast horizon (panel a in figure 2), the error term equals the difference between the realized value of one-year ahead inflation and its consensus forecasted value. For other public signals (panel b in figure 2), the error terms are instead constructed as the residuals from a linear regression of one-year-ahead inflation onto the public signal in question. To make the precision and δ estimates comparable across series, we focus on the longest common sample available (1981Q1–2020Q1) and standardize the variables over this sample. Notice that this contrasts to figure 2, where δ is estimated on the full-sample for each series. Finally, we drop the TIPS from the figure, as it is only available after 2015.

C. *Implications and Discussion*

We conclude this section by discussing auxiliary implications of our calibrated model. The top left-hand panel in Online Appendix F shows the (demeaned) distribution of individual period-*t* forecasts implied by our model. Compared to rational, mean-squared optimal forecasts, the standard deviation of the overconfident forecast distribution is about three times larger. This is because overconfidence causes individuals to put additional weight on private information. Overconfidence in the precision of private information can thus help explain the *a priori* puzzling amount of dispersion in macroeconomic expectations (e.g., Mankiw et al., 2003).

However, importantly, this increase in dispersion does not lead to substantially more imprecise expectations in equilibrium. Online appendix F shows that the standard deviation of errors is only slightly larger in the overconfident case. As a result, forecasters would face difficulty inferring from the accuracy of their own forecasts alone that they were indeed overconfident. The bottom panel in online appendix F illustrates the reason for this similarity: the endogenous public signal (consensus) is substantially more precise in the overconfident case. Because overconfident forecasters put more weight on private information, the endogenous consensus embeds more of the sum of forecasters' private information, the only truly new information that forecasters can learn from each other. In effect, overconfidence counteracts the standard learning externality that exists in markets with public information and which causes agents to attach too little weight to private information (e.g., Amador & Weill, 2010). This connects our results with those of Smith (1982) and others that attempt to find "group optimal explanations" for individual biases.

Finally, a substantial literature in macroeconomics has explored whether noise shocks to public information can explain business cycle fluctuations. Because agents in our model attach more weight to public information than optimal, any such shock also has a heightened effect on individual expectations. This illustrates one potential implication of absolute and relative overconfidence. Other potential implications include: (i) increased trade in financial assets; (ii) "over-shooting" of asset prices in response to public announcements; and (iii) increases in investments into new product lines. We leave these topics, and others, for future research.

VI. Concluding Remarks

In this paper, we have explored the implications of individual professional forecast of macroeconomic variables for popular models of expectation formation. We have demonstrated that the statistical properties of individual forecasts contradict standard versions of noisy rational expectations. In place, we have proposed a simple extension of noisy rational expectations, consistent with the survey evidence, building on two frictions: Forecasters believe that their own private information is not only better than it actually is (absolute overconfidence), but also better than that available to others (relative overconfidence). Combined, these biases entail that forecasters both overreact to private information and misperceive the informativeness of endogenous public signals that aggregate others' private information.

We hope that our paper may serve as a stepping stone for further empirical and theoretical research along similar lines. Our model has illustrated how simple behavioral biases can combine with the endogeneity of public information to create rich patterns of predictability in individual errors. This idea is more general than our particular forecaster application.

REFERENCES

Alicke, M. D., and O. Govorun, "The Better-Than-Average Effect," *The Self in Social Judgment* 1 (2005), 85–106.

Alpert, M., and H. Raiffa, "A Progress Report on the Training of Probability Assessors," in *Judgment under Uncertainty: Heuristics and Biases* (Cambridge: Cambridge University Press, 1982).

Amador, M., and P.-O. Weill, "Learning from Prices: Public Communication and Welfare," *Journal of Political Economy* 118:5 (2010), 866–907. 10.1086/657923

Andrade, P., and H. Le Bihan, "Inattentive professional forecasters," *Journal of Monetary Economics* 60:8 (2013), 967–982. 10.1016/j.jmoneco.2013.08.005

Angeletos, G.-M., and Z. Huo, "Myopia and Anchoring," *American Economic Review* 111:4 (2021), 1166–1200. 10.1257/aer.20191436

Angeletos, G.-M., Z. Huo, and K. A. Sastry, "Imperfect Macroeconomic Expectations: Evidence and Theory," *NBER Macroeconomics Annual* 35:1 (2021), 1–86. 10.1086/712313

Angeletos, G.-M., and A. Pavan, "Efficient Use of Information and Social Value of Information," *Econometrica* 75:4 (2007), 1103–1142. 10.1111/j.1468-0262.2007.00783.x

Armantier, O., G. Topa, W. Van der Klaauw, and B. Zafar, "An Overview of the Survey of Consumer Expectations," *Economic Policy Review* 23:2 (2017), 51–72.

Barberis, N., A. Shleifer, and R. Vishny, "A Model of Investor Sentiment," *Journal of Financial Economics* 49:3 (1988), 307–343. 10.1016/S0304-405X(98)00027-0

Benoît, J.-P., J. Dubra, and D. Moore, "Does the Better-Than-Average Effect Show that People are Overconfident?" *Journal of the European Economic Association* 13:2 (2015), 293–329.

Bordalo, P., N. Gennaioli, and A. Shleifer, "Diagnostic Expectations and Credit Cycles," *The Journal of Finance* 73:1 (2018a), 199–227. 10.1111/jofi.12586

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer, "Overreaction in Macroeconomic Expectations," NBER working paper w24932 (2018b).

——— "Overreaction in Macroeconomic Expectations," *American Economic Review* 110:9 (2020), 2748–2782. 10.1257/aer.20181219

Cameron, A. C., D. L. Miller et al., "Robust Inference with Clustered Data," *Handbook of Empirical Economics and Finance* 106 (2010), 1–28.

Canova, F., "G-7 Inflation Forecasts: Random Walk, Phillips Curve or What Else?" *Macroeconomic Dynamics* 11:1 (2007), 1–30. 10.1017/S136510050705033X

Cecchetti, S. G., "Inflation Indicators and Inflation Policy," *NBER Macroeconomics Annual* 10 (1995), 189–219. 10.1086/654274

Clements, M. P., "Do Macroforecasters Herd?" *Journal of Money, Credit and Banking* 50:2–3 (2018), 265–292. 10.1111/jmcb.12460

Coibion, O., and Y. Gorodnichenko, "What Can Survey Forecasts Tell us About Information Rigidities?" *Journal of Political Economy* 120:1 (2012), 116–159. 10.1086/665662

——— "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review* 105:8 (2015), 2644–78. 10.1257/aer.20110306

Coibion, O., Y. Gorodnichenko, S. Kumar, and J. Ryngaert, "Do You Know that I Know that You Know? Higher-Order Beliefs in Survey Data," *The Quarterly Journal of Economics* 136:3 (2021), 1387–1446. 10.1093/qje/qjab005

Croushore, D. D., "Introducing: The Survey of Professional Forecasters," *Business Review-Federal Reserve Bank of Philadelphia* 6 (1993), 3–15.

——— "The Livingston Survey: Still Useful After all These Years," *Business Review-Federal Reserve Bank of Philadelphia* 2 (1997), 15–27.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam, "Investor Psychology and Security Market Under- and Overreactions," *The Journal of Finance* 53:6 (1998), 1839–1885. 10.1111/0022-1082.00077

De Bondt, W. F. M., and R. Thaler, "Does the Stock Market Overreact?" *The Journal of Finance* 40:3 (1985), 793–805. 10.1111/j.1540-6261.1985.tb05004.x

Dominitz, J., and C. F. Manski, "How Should We Measure Consumer Confidence (Sentiment)? Evidence from the Michigan Survey of Consumers," NBER working paper 9926 (2003).

Ehrbeck, T., and R. Waldmann, "Why are Professional Forecasters Biased? Agency Versus Behavioral Explanations," *The Quarterly Journal of Economics* 111:1 (1996), 21–40. 10.2307/2946656

Einhorn, H. J., "Overconfidence in Judgment," *New Directions for Methodology of Social and Behavioral Science* 4:1 (1980), 1–16.

Evans, G. W., and S. Honkapohja, *Learning and Expectations in Macroeconomics* (Princeton, NJ: Princeton University Press, 2012).

Eyster, E., M. Rabin, and D. Vayanos, "Financial Markets Where Traders Neglect the Informational Content of Prices," *The Journal of Finance* 74:1 (2019), 371–399. 10.1111/jofi.12729

Garcia, J. A., "An Introduction to the ECB's Survey of Professional Forecasters," ECB occasional paper 8 (2003).

Giordani, P., and P. Söderlind, "Inflation Forecast Uncertainty," *European Economic Review* 47:6 (2003), 1037–1059. 10.1016/S0014-2921(02)00236-2

Goldstein, N., "Tracking Inattention," mimeo (2021).

Graham, J. R., "Herding Among Investment Newsletters: Theory and Evidence," *The Journal of Finance* 54:1 (1999), 237–268. 10.1111/0022-1082.00103

Griffin, D., and A. Tversky, "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology* 24:3 (1992), 411–435. 10.1016/0010-0285(92)90013-R

Hirshleifer, D., S. S. Lim, and S. H. Teoh, "Limited Investor Attention and Stock Market Misreactions to Accounting Information," *The Review of Asset Pricing Studies* 1:1 (2011), 35–73. 10.1093/rapstu/rar002

Kohlhas, A. N., and A. Walther, "Asymmetric Attention," *American Economic Review* 111:9 (2021), 2879–2925. 10.1257/aer.20191432

Köhler, M., and S. Schmidt, "ZEW Finanzmarkreport: Kurz Info," mimeo (2021).

Lamont, O. A., "Macroeconomic Forecasts and Microeconomic Forecasters," *Journal of Economic Behavior & Organization* 48:3 (2002), 265–280.

Larrick, R. P., K. A. Burson, and J. B. Soll, "Social Comparison and Confidence: When Thinking You are Better Than Average Predicts Overconfidence (and When it Does Not)," *Organizational Behavior and Human Decision Processes* 102:1 (2007), 76–94. 10.1016/j.obhdp.2006.10.002

Laster, D., P. Bennett, and I. S. Geoum, "Rational Bias in Macroeconomic Forecasts," *The Quarterly Journal of Economics* 114:1 (1999), 293–318. 10.1162/003355399555918

Lorenzoni, G., "A Theory of Demand Shocks," *American Economic Review* 99:5 (December 2009), 2050–2084. 10.1257/aer.99.5.2050

Maćkowiak, B., and M. Wiederholt, "Optimal Sticky Prices Under Rational Inattention," *The American Economic Review* 99:3 (2009), 769–803.

Mankiw, N. G., and R. Reis, "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics* 117:4 (2002), 1295–1328. 10.1162/003355302320935034

Mankiw, N. G., R. Reis, and J. Wolfers, "Disagreement About Inflation Expectations," *NBER Macroeconomics Annual* 18 (2003), 209–248. 10.1086/ma.18.3585256

Marinovic, I., M. Ottaviani, and P. Sorensen, "Forecasters Objectives and Strategies" (pp. 690–720), in *Handbook of Economic Forecasting*, Vol. 2 (Amsterdam: Elsevier, 2013). 10.1016/B978-0-444-62731-5.00012-9

Moore, D. A., and P. J. Healy, "The Trouble with Overconfidence," *Psychological Review* 115:2 (2008), 502. 10.1037/0033-295X.115.2.502

Muth, J. F., "Rational Expectations and the Theory of Price Movements," *Econometrica: Journal of the Econometric Society* 29 (1961), 315–335. 10.2307/1909635

Ottaviani, M., and P. N. Sørensen, "The Strategy of Professional Forecasting," *Journal of Financial Economics* 81:2 (2006), 441–466. 10.1016/j.jfineco.2005.08.002

Scharfstein, D. S., and J. C. Stein, "Herd Behavior and Investment," *The American Economic Review* (1990), 465–479.

Sims, C. A., "Implications of Rational Inattention," *Journal of Monetary Economics* 50:3 (2003), 665–690. 10.1016/S0304-3932(03)00029-1

Smith, J. M., *Evolution and the Theory of Games* (Cambridge: Cambridge University Press, 1982).

Soll, J. B., and J. Klayman, "Overconfidence in Interval Estimates," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 30 (2004), 299. 10.1037/0278-7393.30.2.299

Stock, J. H., and M. W. Watson, "Phillips Curve Inflation Forecasts" (pp. 101–186), in Jeff Fuhrer, Yolanda K. Kodrzycki, Jane Sneddon Little, and Giovanni P. Olivei (eds.), *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective* (Cambridge, MA: MIT Press, 2008).

Veldkamp, L. L., *Information Choice in Macroeconomics and Finance* (Princeton, NJ: Princeton University Press, 2011).

Woodford, M., "Imperfect Common Knowledge and the Effects of Monetary Policy," NBER working paper 8673 (December 2002).