Forecaster (Mis-)Behavior*

Tobias Broer† Alexandre N. Kohlhas‡

November, 2018

Abstract

We document how US professional forecasters overrespond to private and public information in their inflation forecasts. We show that such overresponses are inconsistent with rational expectation formation with noisy information, common agency-based models of forecaster behaviour, as well as several behavioural alternatives. In place, we propose a simple model of naïve overconfidence, consistent with the stylized facts. Unlike rational forecasters, naïve forecasters are overconfident in their private information and believe it to be superior to that held by others. We show how such naivety causes forecasters to misinterpret the information content of market-generated public outcomes. As a result, forecasters not only overrespond to private information but also to observed public information. The model sheds light on the importance of the misinterpretation of others' behaviour for micro-consistent models of expectation formation.

JEL codes: C53, D83, D84 *Keywords:* Expectations, forecasters, information

1 Introduction

Expectations about tomorrow matter for choices today. The design of good macroeconomic policies therefore depends on the way households and firms make predictions about the future. Because individual expectations are typically unobserved, however, it is difficult to discriminate between different models of expectation formation. One exception are professional forecasters, who regularly publish predictions about major macroeconomic and financial variables.

*First draft: March, 2018. We thank Zhen Huo, Per Krusell, Kristoffer Nimark, Kurt Mitman, Torsten Persson, Ansgar Walther, and audiences at several seminars for comments. We also thank John Kramer for outstanding research assistance. Support from Ragnar Soderbergs Stiftelsen is gratefully acknowledged.

†Address: Institute for International Economic Studies, SE-106 91 Stockholm, Sweden.
Email: `tobias.broer@iies.su.se`; website: `http://perseus.iies.su.se/tbroe`.

‡Address: Institute for International Economic Studies, SE-106 91 Stockholm, Sweden.
Email: `alexandre.kohlhas@iies.su.se`; website: `https://alexandrekohlhas.com`.

Using such forecasts to gauge economic expectations more generally is nevertheless problematic for two reasons: First, professional forecasters are likely to differ from other economic agents in their personal characteristics and information about the state of the economy. And second, their forecasts may be determined by strategic behavior and thus not accurately reflect the mean of their posterior distribution of the variable of interest.

Recently, survey data on professional forecasts have been found to support the main hypotheses of a growing literature on rational expectation formation with imperfect information. Specifically, Coibion and Gorodnichenko (2015) and Dovern *et al.* (2015) find that on *average* professional forecasts *underreact* to news, in the sense that forecast revisions are too small.[1] This underresponse is in turn consistent with rational behavior by forecasters whose noisy private information optimally dampens individual forecast revisions.

Using the US Survey of Professional forecasters (SPF), we find that *individual* SPF inflation forecasts indeed *overreact* to information, in the sense that their forecast revisions are too large. Importantly, this holds in response to private information, and in response to public information in the form of previous consensus forecasts. Specifically, positive individual forecast revisions are associated with over-predictions of the variable of interest, and vice versa for negative forecast revisions. And a higher average, or 'consensus', forecast is associated with a larger over-prediction in later forecasts.

To explain these facts, we consider two models of expectation formation that have been found to imply overreaction to news. The first is a model of strategic diversification as in Laster *et al.* (1999), Marinovic *et al.* (2013), and Marinovic *et al.* (2013) where forecasters are paid for correctly predicting outcomes but payoffs are higher when there are only few other correct forecasts. Such a contest motive naturally gives forecasters incentives to distort their forecasts away from the mode of the distribution of forecasts, towards their private information. We compare this 'strategic' model of forecaster behavior to a model of overconfidence: when forecasters are overconfident about the accuracy of their own private information relative to that of others, they naturally overreact to private signals.

We show how standard departures from the simple noisy rational expectations framework may explain overreaction to private information, but cannot explain overreaction to public information. The reason is that w

This paper is related to the literature on expectation formation, as well as to studies of the behavior and incentives of professional forecasters. It is most related to Coibion and Gorodnichenko (2015), who show that average forecasts of several key economic variables in both the US SPF and other forecaster surveys exhibit under-responses to news, in the sense that forecast revisions are positively correlated with forecast errors. Their study builds on a

[1] Bordalo *et al.* (2018a) and Fuhrer (2015) provide related evidence.

number of previous tests of the hypothesis of limited information rational expectations (LIRE), starting with Mankiw and Reiss (2005) who find support for a sticky-information model in both consumer expectations and SPF forecasts.[2]. Our contribution is to study both individual and average SPF forecasts within a consistent framework.[3] We show that forecast revisions are negatively correlated with forecast errors at the individual level, and that there is a significant negative correlation between consensus forecasts and individual forecast errors. We argue that this can be interpreted as overreaction by forecasters to both private and public information.

Although forecaster information is sometimes acknowledged to be an upper bound of that held by the population at large Andrade and Le Bihan (2013), most studies that use forecaster surveys to test the LIRE hypothesis abstract from the particular characteristics that distinguish professional forecasters from the rest of the population. This has attracted criticism (for example by Lamont, 2002) and given rise to a literature that looks at the incentives to distort forecasts away from true conditional expectations of future variables. For example, forecasters might try to mimic their more able colleagues (as in Ehrbeck and Waldmann, 1996), or, more generally, change their forecasts as a function of those they anticipate from their colleagues (Lamont, 2002). We focus on one particular incentive distortion that has been found to imply overreaction in static environments and is sometimes called the 'contest' hypothesis (Laster *et al.*, 1999, Ottaviani and Sørensen, 2006, and Marinovic *et al.*, 2013), whereby the payoff from a correct forecast is higher when only few competitors make the same forecast. In a static context with a common prior, \citet{ottaviani2006strategy} show how this makes forecasters optimally distort their forecasts away from their posterior mean and towards their idiosyncratic signal. We extend their framework to a dynamic environment with repeated private signals and public information in the form of consensus forecasts. We show that the contest motive implies overreaction to private signals also in that context. Our contribution is to show that the rationality at the core of the model limits the overreaction to public signals, inconsistent with the substantial overresponse that we find in the data.

The literature on forecaster behavior that accounts for their particular strategic incentives typically maintains the assumption of rational expectation formation. Research on alternatives to the LIRE assumption has been very active over recent years, but typically looks at anomalies in financial data, rather than in macroeconomic forecasts. One exception is Bordalo *et al.* (2018b), who document that survey forecasts of credit spreads feature periods of excessive optimism and, more generally, predictable errors. Based on Kahneman and Tversky's (1973) representative heuristic they explain this by 'diagnostic' expectation formation

[2]See Coibion and Gorodnichenko (2015) for a detailed survey of that literature

[3]Dovern *et al.* (2015) look at revisions of individual GDP forecasts for a broad set of countries. They argue that these are less supportive of sticky information rational expectation models (Reiss, 2006) relative to their noisy information counterpart (Woodford, 2001)

Gennaioli and Shleifer (2010), which exaggerates the probability of future states whose likelihood has increased by recent news. A particular well-studied departure from LIRE is investor overconfidence in the precision of their own knowledge. For example, Ben-David et al (2003), show how executives are on average overconfident in their ability to predict stock returns, and that this overconfidence correlates with corporate policies (as also shown by Malmendier et al, 2005). We compare the LIRE and contest hypotheses to a simple version of overconfidence, which makes forecasters believe that their signals are more precise than those of others. We show how this implies, perhaps surprisingly, overreaction to both private and public signals as well as over-correction of previous deviations from consensus that are not too far from those observed in the data.

Section 2 presents a standard framework of expectation formation with limited noisy information, and shows how this implies under-reaction of average expectations to news. Section 3 uses US SPF data to show how there is indeed under-reaction of average forecasts to average news, as found in Coibion and Gorodnichenko (2015). Individual forecasts of both inflation measures and real gross national output, in contrast, overreact to idiosyncratic news. They also overreact to public news contained in consensus forecasts. Section 4 presents two alternative models and derives their theoretical properties. Section 5 estimates the two models using a simulated method of moment procedure and compares their key predictions to the data. Section 6 concludes. An appendix contains further empirical results and all proofs.

2 A Baseline Model

We start with a simple model of rational forecasting with dispersed information, to fix ideas about which hypotheses we later on test in the survey data.

There is a continuum of measure one of forecasters, indexed by the unit interval $i \in [0, 1]$. Each forecaster i minimizes the mean-squared error of his forecast f_i of the random variable θ. of which the forecaster has the prior belief $\theta \sim N(\mu_i, \tau_\theta^{-1})$. All forecasters observe two types of information. Their own private information, summarized by the *private signal*

$$x_i = \theta + \epsilon_i, \tag{2.1}$$

where the noise terms $\epsilon_i \sim \mathcal{N}(0, \tau_x^{-1})$ are independent of θ and $\mathbb{E}[\epsilon_i \epsilon_j] = 0$ for all $j \neq i$. The private signal of one forecaster is not observed by any other forecaster. This is the sense in which the information encoded in x_i is private. In addition to their private information, all forecasters observe the realization of the *public signal*

$$y = \theta + \xi, \tag{2.2}$$

where $\xi \sim \mathcal{N}(0, \tau_y^{-1})$ is independent of θ and ϵ_i for all $i \in [0,1]$. The signal y is public in that its realization is common knowledge among all forecaster.

Several important implications of rational expectations and dispersed information can be studied within this framework.[4] We focus on three implications of $f_i = \mathbb{E}[\theta \mid \mu_i, x_i, y]$ below.[5]

Average Forecasts (Implication 1): Consider the problem faced by individual i. Based on both private and public information, his optimal forecast of θ is

$$f_i = \mathbb{E}[\theta \mid \mu_i, x_i, y] = \mu_i + k\left(\mathbb{E}[\theta \mid x_i, y] - \mu_i\right), \tag{2.3}$$

where $k = \frac{\tau_x + \tau_y}{\tau_\theta + \tau_x + \tau_y} \in (0,1)$ denotes the combined weight on new information, x_i and y. Averaging (2.3) across all $i \in [0,1]$ and rearranging we arrive at

$$\theta - f = \frac{1-k}{k}(f - \mu) - c\xi,$$

where $f = \int_0^1 f_i di$ denotes the average forecast, $\mu = \int_0^1 \mu_i di$ the average prior expectation, $c = \frac{\tau_y}{\tau_x + \tau_y}$ and $\int_0^1 \epsilon_i di = 0$ since the noise in the forecasters' private signals cancels on average.

Estimates of the slope b in a regression of mean forecast errors $\theta - f$ on mean forecast revisions $f - \mu$,

$$\theta - f = a + b(f - \mu) + \nu, \tag{2.4}$$

where α and ν denote constant and error terms, respectively, will therefore be positive and measure the extent of information frictions ($k < 1$). Estimates of $b > 0$ will measure the extent to which forecasters' information (x_i and y) accurately reflect the forecast variable θ.

Because of the presence of private information, individuals, on average, underrespond to new information in (2.4), as captured by the average forecast revision. Each forecaster down-weighs his own information to account for its noisiness $k < 1$. But since the private noise terms cancel on average, this down-weighing of information leads to an underresponse to the average new information observed, and hence to a positive correlation between the mean-forecast error, on the one hand, and the mean-forecast revision, on the other hand.

Indeed, as shown by Coibion and Gorodnichenko (2015), the correlation of errors and revi-

[4] Apart from those discussed below, several broad aspects of survey forecasts are clearly consistent with noisy rational expectations. First, survey forecasts are dispersed and differ across forecasters (Zarnowitz, 1985). Second, forecasts are often smoother, with lower volatility, than the variable that is being forecasted (Ottaviani and Sørensen, 2006). In fact, one of Muth's (1961) aims in proposing the rational expectations hypothesis was to explain these two stylized facts observed using survey data (p. 361 in Muth, 1961).

[5] As discussed in e.g. Bhattacharya and Pfleiderer (1985) conditional expectations do not only correspond to optimal predictors for mean-squared error loss functions but indeed for any symmetric loss function.

sions is generically positive for average forecasts across a wide-variety of information structures when forecasters are in possession of private information. The mean forecast, in essence, encodes more precise information than used by individual forecasters because of the presence of private information. This, in turn, causes the average forecast revision to underrespond to average new information. The averaging of private informations also renders the average forecast more precise (Zarnowitz and Lambros, 1987).

Individual Forecasts (Implication 2 and 3): A well-known consequence of conditional expectation forecasts is that individual forecast errors $\theta - f_i$ are uncorrelated with convex combinations of the elements in forecasters' information set (Granger, 1969).

Let z denote any linear combination of x_i, y and μ_i. Then,

$$\mathbb{E}\left[\theta - f_i \mid z\right] = \mathbb{E}\left[\theta \mid z\right] - \mathbb{E}\left[\mathbb{E}\left[\theta \mid x_i, y, \mu_i\right] \mid z\right] = 0, \tag{2.5}$$

since $\mathbb{E}\left[\mathbb{E}\left[\theta \mid x_i, y, \mu_i\right] \mid z\right] = \mathbb{E}\left[\theta \mid z\right]$ by the *Law of Iterated Expectations.*

Equation (2.5) has two important consequences. First, estimates of the slope in (2.4) at the individual level should be zero. That is, estimates of β in the regression

$$\theta - f_i = \alpha + \beta\left(f_i - \mu_i\right) + \nu_i, \tag{2.6}$$

should be undistinguishable from nil, in contrast to at the average level where the slope coefficient is positive ($b > 0$). An individual's forecast error $\theta - f_i$ cannot be correlated with his own forecast revision $f_i - \mu_i$, itself a combination of x, y and μ_i. If it was, the forecaster could exploit this correlation to improve his forecast, contradicting the assumption that his forecast was the conditional expectation to start with. And second, by similar logic, (2.5) also dictates that individual forecast errors should be uncorrelated with any public information y that forecast revisions are in part based on. The estimate of δ in the regression

$$\theta - f_i = \alpha + \delta y + \nu_i. \tag{2.7}$$

should thus also equal zero, since any non-zero coefficient would contradict the assumption that forecast correspond to conditional expectations.

We summarize these implications of noisy rational expectations in Proposition 1

Proposition 1. *With rational expectations and private information, individual forecast revisions and public information do not predict individual forecast errors, $\beta = 0$ and $\delta = 0$. Mean forecast errors are, however, positively correlated with mean forecast revisions, $b > 0$.*

Because of the potential for heterogenous priors $\mu_i \neq \mu_j$ for $i \neq j$, Proposition 1 covers the important case in which θ itself evolves dynamically across time, in accordance with for example an $AR(1)$. Forecasts in (2.3) then correspond to those that arise from the Kalman Filter (Anderson and Moore, 1979) when new information is observed in each period.[6]

3 Empirical Evidence

In this section, we compare the implications of noisy rational expectations, listed in Proposition 1, to key features of US inflation forecasts. We document how professional forecasters on average underrespond to new information ($b > 0$). We then show how the same forecasters at the individual level in contrast overrespond to their own forecast revision ($\beta < 0$) as well as to observed public information ($\delta < 0$). Last, we show how these results carry over to other forecast variables, such as output, alternative forecast periods, and datasets.

Data: We focus on forecasts of US inflation 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 after. We study SPF forecasts of the year-on-year percentage change in the GNP/GDP deflator, for which the survey includes consistent forecasts over 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 with those from previous studies, which focus disproportionally on inflation expectations. And last, because data are available for a substantially longer time-span for inflation forecasts in the US than for other variables, such as output, or other countries. Throughout, we consider first-release realizations of inflation to most accurately capture the precise definition of the variable being forecasted (Croushore and Stark, 2001). Importantly for our purposes, although the precise schedule over the quarter has changed over time, the administrators of the SPF have consistently and publicly published the survey results well before sending out the following questionnaire to their respondents.[8] We therefore assume that the information

[6] Furthermore, because the last two implications in Proposition 1 ($\beta = 0$ and $\delta = 0$) derive from the Law of Iterated Expectations these extend to economies in which shocks are non-normal. The first implication, $b = 0$, carries directly over to other affine prior-likelihood combinations, such as beta-binomial, gamma-poisson, and when observations are negative binomial, gamma, or exponential with natural conjugate priors (Ericson, 1969).

[7] The SPF is the oldest quarterly survey of individual, macroeconomic forecasts in the US, dating back to 1968. The SPF was initiated under the leadership of Arnold Zarnowitz at the American Statistical Association and the National Bureau of Economic Research, which is why it is also still occasionally referred to as the the ASA-NBER Quarterly Economic Outlook Survey (Croushore, 1993).

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

set of respondents includes the consensus (or average) forecast of the previous quarter.[9]

Average Forecasts: We first study the properties of average inflation forecasts. We denote individual i's forecast made in period t of the variable π in period $t+h$ as $f_{it|t+h}$ and calculate the average as $f_{t|t+h} = \frac{1}{N_t}\sum_{i=1,..N_t} f_{it|t+h}$, where N_t denotes the number of forecasters in period t. We then estimate the following regression equation

$$\pi_{t+h} - f_{t|t+h} = a + b\left(f_{t|t+h} - f_{t|t+h-1}\right) + \nu_t. \tag{3.1}$$

We thus estimate (2.4) when θ corresponds to inflation h periods from now.[10]

Table I presents the results for one-year ahead inflation forecasts, $h = 4$. Average forecasts revisions are positively correlated with average forecast errors, $\hat{b} > 0$. This effect is strong and highly significant, in accordance with the results in Coibion and Gorodnichenko (2015), Andrade and Le Bihan (2013), Dovern *et al.* (2015), and Bordalo *et al.* (2018a). Consistent with noisy rational expectations, professional forecasters on average underrespond to the new information that they observe between two forecast rounds, leading to a positive correlation between the average forecast error, on the one hand, and the average forecast revision, on the other hand. Furthermore, consistent with the averaging out of noise in private information, Table I shows how the average forecast is also more precise than that of any individual forecast. This corroborates the empirical findings of Zarnowitz and Lambros (1987).

Individual Forecasts: We now turn to our main empirical results, describing the statistical properties of individual inflation forecasts. In particular, we test the implication of rational expectations, detailed in Proposition 1, that forecast errors are orthogonal to any information, be it public or private, available at the time of the forecast. Figure 3 shows that this implication is *prima facie* not borne out by the data. The means of individual forecast errors are negatively associated with the means of individual forecast revisions (left panel) as well as consensus forecasts for the same period from the previous wave of the survey (right panel).

To test this implication more formally, we first estimate the equivalent of (2.6) at the individual level, using the benchmark specification

$$\pi_{t+h} - f_{it|t+h} = \alpha_i + \beta\left(f_{it|t+h} - f_{it|t+h-1}\right) + \nu_{it}, \tag{3.2}$$

[9]The number of respondents (professional forecasters from financial institutions, large industrial firms, and independent forecasting enterprises) fell from over 80 to less than 20 in 1990, when the Federal Reserve Bank of Philadelphia took over the administration, and has fluctuated between 40 and 60 since then. See `https://www.philadelphiafed.org/-/media/research-and-data/real-time-center/survey-of-professional-forecasters/spf-documentation.pdf?la=en` for more details.

[10]Notice that because of the presence of public information least-squares estimates of b in (2.4) and (3.1) are downward bias. As argued in Coibion and Gorodnichenko (2015), such downward bias, however, still entails that statistically significant findings of $b > 0$ are valid since estimates will understate the positive association.

Figure 1: Forecast Errors from the Survey of Professional Forecasters

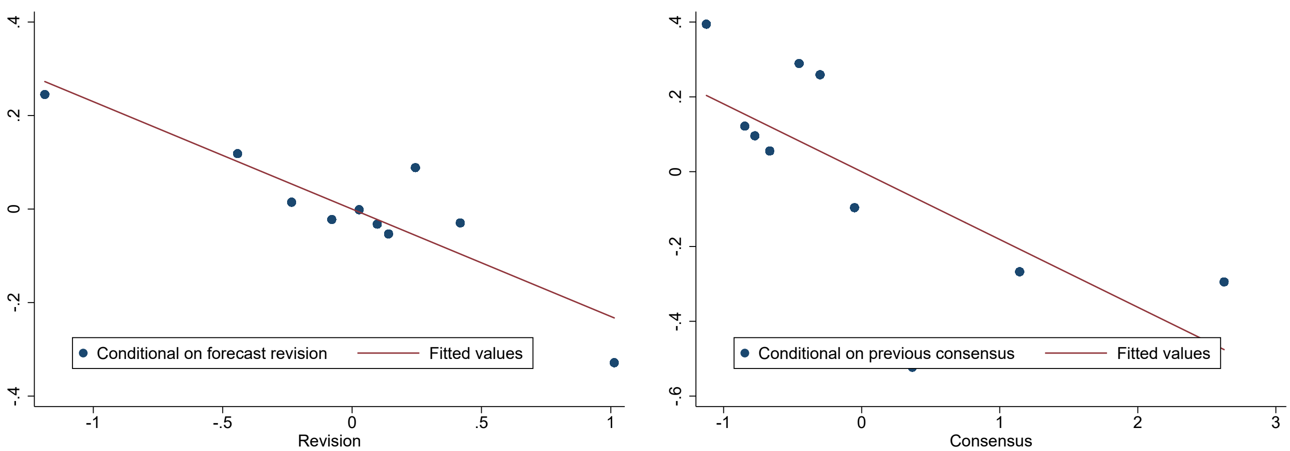

The figure depict the means of forecast errors of inflation taken within deciles of the distributions of forecast revisions (*left panel*) and consensus forecasts for the same period in the previous wave of the SPF (*right panel*). All variables are demeaned by subtracting their overall mean during the sample period (1970Q1-2016Q4).

Table I: Estimates from the Survey of Professional Forecasters

	Average Forecasts	*Individual Forecasts*		
	Forecast Error	*Forecast Error*	*Forecast Error*	*Forecast Error*
Forecast Revision	1.273***	-0.195***	–	-0.199***
	(0.276)	(0.0432)		(0.0432)
Previous Consensus	–	–	-0.169***	-0.166***
			(0.0343)	(0.0446)
Constant	-0.0750	-0.0168***	0.604***	0.538***
	(0.0732)	(0.0000626)	(0.120)	(0.149)
Sample	01/70:10/16	01/70:10/16	01/70:10/16	01/70:10/16
Obs	182	5016	6722	5016
R^2	0.255	0.020	0.017	0.037

(i) Column one presents estimates of (3.1); column two and three (3.2) and (3.3).

(ii) HAC standard errors used. * p<0.10, ** p<0.05, *** p<0.01.

where α_i denotes a forecaster specific fixed-effect.[11] Equation (3.2) thus corresponds to (2.6) when θ is once more set equal to h-period ahead inflation, π_{t+h}.

Our second test instead estimates a version of (2.7) by using a particular piece of public information that is presumably salient to professional forecasters – namely the consensus forecast from the previous wave of the survey:

$$\pi_{t+h} - f_{it|t+h} = \alpha_i + \delta f_{t-1|t+h} + \nu_{it}. \tag{3.3}$$

As argued above, and more 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 immediate competitors. Consensus estimates should therefore provide a conservative benchmark against which to test the orthogonality of individual forecast errors to public information implied by noisy rational expectations.[12]

Table I presents estimates of (3.2) and (3.3) for one-year ahead inflation forecasts, $h = 4$. Proposition 1 showed that under dispersed information and rational expectations, we should expect estimates for the slope coefficients β and δ in both regressions close to zero. By contrast, the estimated β and δ in Table I are significantly negative and numerically large, inconsistent with the noisy rational expectations. The negative estimated value of β implies that positive individual forecast revisions are associated with negative forecast errors. Thus, forecasters on average revise their forecasts by too much, and hence overreact to the information recevied between subsequent survey rounds. Importantly, the negative estimate of δ implies that such overresponses also hold in reaction to a particular element of public information: Individual forecast errors are on average more negative not only when individual forecast revisions are more positive but also when the previous consensus forecast is higher. These overresponses are corroborated in the final column of Table I, where we report the coefficient estimates of the multivariate regression, controlling both for individual forecast revisions and previous consensus realizations. Forecasters indeed overrespond both to the public information embodied in consensus forecasts and to the remaining information that they receive between forecast rounds conditioning on consensus.

Robustness: These patterns documented in the SPF are remarkably stable when we consider other macroeconomic variables, forecast periods, forecast horizons, and countries. Appendix A summarizes the results from alternative estimates of (3.1), (3.2), and (3.3).

First, to complement our benchmark results using GNP/GDP inflation forecasts, we also consider forecasts of CPI inflation and real GNP/GDP growth from the SPF. The estimated

[11] The removal of the fixed-effect in our benchmark specification does not affect the insights from Table I.

[12] Pesaran and Weale (2006) and Fuhrer (2015) document how the orthogonality restriction of individual forecast errors to public information also fails for other variables than consensus.

coefficients for b, β, and δ all have the same sign as our benchmark results, and are all statistically significant with the exception of the CPI estimate of β. Similar results holds when we restrict the sample to the period after the Federal Reserve Bank of Philadelphia took over the administration of the SPF in 1993 and substantially increased its coverage. We also document that similar results hold at a semi-annual forecast horizon, where $h = 2$.

Second, we extend beyond the United States and consider professional forecasts of inflation for another geographic area, the Euro Area, as collected by the *ECB's Survey of Professional Forecasts*. We once more find estimates of the same sign and similar magnitudes to those from the US SPF, albeit the estimate on forecast revisions looses its statistical significance. We attribute this result to the low power of any test implied by the short estimation sample.[13]

Apart from these additional estimates, our confidence in the overresponse to individual forecast revisions ($\beta < 0$) is also corroborated by contemporaneous work. In their paper, Bordalo *et al.* (2018a) document a significant negative correlation between individual forecast errors and individual forecast revisions across many different economic variables in a wide variety of datasets that much extends beyond those considered in Table I and Appendix A. They do not, however, study the relationship between forecast errors and public information.

Taken together, the data thus strongly suggest that underresponses ($b > 0$) and overresponses ($\beta < 0$ and $\delta < 0$) occur simultaneously. Forecasters, on average, underrespond to the new information that they receive between two periods, consistent with noisy rational expectations. But, at the individual level, forecasters instead overrespond to the news that they receive. As we have argued in the introduction, and will show more formally below, models based purely on noisy information and rational forecaster incentives have a hard time to reconcile this coincidence of over- and underresponsive forecasts at the individual level. The next section makes this more explicit using the basic framework from Section 2.

4 A Taxonomy of Explanations

A variety of explanations are consistent with the simultaneous under- and overresponse of forecast revisions at the average ($b > 0$) and individual level ($\beta < 0$), respectively. On balance, these explanations can be separated into two classes: First, agency-based explanations in which the assumption of mean-squared error loss is replaced while the assumption of rational information use is maintained. And second, behavioral explanations, such as overconfidence (or representativeness), that keep the assumption of mean-squared error loss but allow for non-rational uses of information. In this section, we show how several of the most prominent of such explanations are, nevertheless, inconsistent with the overresponse to public information

[13]The ECB's Survey of Professional Forecasts starts in 2000.

($\delta < 0$) that we further documented in the previous section.

Private Information Responses: Consider once more individual i's forecast of θ from Section 2, $f_i = \mathbb{E}[\theta \mid \mu_i, x_i, y]$. We can alternatively write this forecast as

$$f_i = (1 - w_\star)\,\mathbb{E}[\theta \mid \mu_i, y] + w_\star x_i, \tag{4.1}$$

where $w_\star = \frac{\tau_x}{\tau_\theta + \tau_x + \tau_y}$ denotes the mean-squared optimal weight on private information. But suppose now that this weight differs from $w_\star$, so that f_i is instead characterized by

$$f_i = (1 - w)\,\mathbb{E}[\theta \mid \mu_i, y] + w x_i, \quad w > w_\star. \tag{4.2}$$

It then follows from (4.3) and (4.4) that[14]

$$\mathbb{C}ov\left[\theta - f_i, f_i - \mu_i\right] = -w\left(w - w_\star\right)\mathbb{E}\left[x_i - \mathbb{E}[x_i \mid \mu_i, y]\right]^2 < 0$$

when $w > w_\star$. When individuals attach more weight to their private information than optimal, individuals will, on average, overrespond to the information that they receive between two periods. This, in turn, leads to a negative relationship between individual forecast errors $\theta - f_i$ and forecast revisions $f_i - \mu$ ($\beta < 0$ when $w > w_\star$), precisely as in the data.

Furthermore, it is straightforward to show that as long as $w < 1$, so that individuals respond less to their private information than the optimal response to the average private signal $\int_0^1 x_i di = \theta$, this negative relationship coincides with an average underresponse to the average forecast revision; that is with estimates of b that are strictly positive ($b > 0$).

Public Information Responses: While an increased use of private information can lead to the overresponse to forecast revisions at the individual level ($\beta < 0$) as well as underresponse at the average level ($b > 0$), it nevertheless does not *per se* lead to an overresponse to public information. In fact, despite this altered use of private information, individual forecast errors will still be uncorrelated with public information.

Consider the forecast error based on (4.2)

$$\theta - f_i = \theta - w x_i - (1 - w)\,\mathbb{E}[\theta \mid \mu_i, y] \tag{4.3}$$

Taking conditional expectations based upon public information y then directly shows that

$$\mathbb{E}\left[\theta - f_i \mid y\right] = (1 - w)\left(\mathbb{E}\left[\theta \mid y\right] - \mathbb{E}\left\{\mathbb{E}\left[\theta \mid \mu_i, y\right] \mid y\right\}\right) = 0, \tag{4.4}$$

[14] Add derivation in footnote and show the below statement also

where the last equity follows from the Law of Iterated Expectations. As a result, despite that individual forecasters potentially misuse their private information their forecast errors will still be uncorrelated with public information, $\delta = 0$. The altered use of private information (relative to mean-squared optimal) does not *per se* entail a relationship between individual forecast errors and, for example, consensus, as otherwise documented in Table I.

When one considers estimates of forecast errors onto public information, one only considers whether those sources of information are employed to minimize forecast errors -- and not whether all sources of information in general are accurately employed. Although forecasters with $w \neq w_\star$ do not optimally use their private information to minimize their forecast errors, conditional on this misuse they still utilize public information efficiently. That is why $\mathbb{E}[\theta \mid \mu_i, y]$ appears in (4.2). This, in turn, leads to estimates of δ that are equal to zero.

We summarize the results from this subsection and the previous in Proposition 2.

Proposition 2. *Consider a forecast by individual $i \in [0,1]$ of the form*

$$f_i = (1-w)\,\mathbb{E}[\theta \mid \mu_i, y] + w x_i, \quad 1 > w > w_\star \tag{4.5}$$

Then, $b > 0$ in (2.4), $\beta < 0$ in (2.6), but $\delta = 0$ in (2.7).

Alternative Explanations: A variety of models of forecaster behavior fall within the class described by Proposition 2. These models are thus consistent with a simultaneous underresponse to news at the average level (and thus with $b > 0$), an overresponse at the individual level (implying $\beta < 0$), but inconsistent with the documented overresponse to public information (as evidenced by $\delta < 0$).

Strategic Diversification: Laster *et al.* (1999), Ottaviani and Sørensen (2006) and Marinovic *et al.* (2013) propose agency-based models of professional forecasters that fall within the class described in Proposition 2. In their model, the market for professional forecasters corresponds to a winner-takes-all competition, where only the most accurate forecast is rewarded by a fixed payoff that is split equally among all winners. As a consequence, the equilibrium distribution of forecast becomes an important determinant of individual forecasts. In a symmetric equilibrium all forecasters choose to over-emphasize private information and follow (4.5) with $w > w_\star$. Consider an individual forecaster who sets $w = w_\star$. Increasing his weight on private information ($w > w_\star$) leaves his probability of winning the contest approximately unchanged (as his posterior is flat at the conditional expectation), but strictly reduces the mass of other forecasters that makes the same forecast. In equilibrium, all forecasters therefore choose to set $w > w_\star$. In accordance with Proposition 2, this entails a simultaneous under- and overresponse to forecast revisions ($b > 0$ and $\beta < 0$). But since public information does not diversify forecasts away from those of others, it is still the case that $\delta = 0$.

Reputational Considerations: Ehrbeck and Waldmann (1996) propose an alternative model of forecaster behavior that builds on the reputational considerations faced by professional forecasters. In their model, one set of professional forecasters has access to more precise private information than another. Since forecasters are rewarded by clients by the perceived accuracy of their forecasts, the set of forecasters that receive less precise private information overrespond to their private information in an attempt to mimic their more informed counterparts. **Appendix B** and Ehrbeck and Waldmann (1996) show that these less able forecasters in equilibrium follow (4.5), where $w > w_\star$, while their more informed counterparts in simply set $w = w_\star$. Reputational considerations are thus simultaneously consistent with the documented underresponse to news at the average level ($b > 0$) as well as the overresponse at the individual level ($\beta < 0$). But since forecasters do not differ in their access to public information, all forecasters still correctly use realizations of public information ($\delta = 0$).

Behavioral Overprecision: A considerable literature in psychology has documented how people tend to over-emphasize their own information (see, for instance, Kahneman and Tversky, 1973). As discussed in, for example, Cutler *et al.* (1990), and more recently in Bordalo *et al.* (2018a) and Bordalo *et al.* (2018b) such overconfidence could also provide the basis for the apparent overresponses to news observed in financial market forecast data, over and above that documented in the macroeconomic forecast data in Section 3. Overconfident forecasters believe the precision of their private information to be higher than it actually is ($\tau_x' > \tau_x$). Their forecasts thus follow (4.5) but where $w \in (w_\star, 1)$. However, as argued above, such overconfidence in private information would *in and of itself* not result in overresponses to public information since it does not also entail an amended use of public information.

Other Models of Forecaster Behavior: We conclude this section by mentioning that several other, prominent explanations, either agency-based or behavioral, have been proposed to describe the behavior of professional forecasters. Graham (1999), Lamont (2002), Ottaviani and Sørensen (2006) describe different models in which forecasters all have an incentive to heard, as in Scharfstein and Stein (1990).[15] Ehrbeck and Waldmann (1996), by contrast, describe an alternative model of reputational concerns, where able forecasters have more precise prior information. Last, Daniel *et al.* (1998) detail a model in which security analysts for behavioral reasons underrespond to their own private news. All of these explanations fall within the class described in Proposition 2, but where $w \in (0, w_\star)$. As a result, these models cannot explain the overresponses to individual forecast revisions described in Section 3.

[15] See also, for example, Croushore (1997) and Welch (2000).

5 Misperceptions of Public Information

The previous section showed how a variety of prominent, proposed models of forecaster behavior are inconsistent with the documented overresponse to public information. ($\delta < 0$) This is despite the capacity for these models to match the simultaneous under- and overresponse of forecast revisions at the average ($b > 0$) and individual level ($\beta < 0$), respectively. In this section, we characterize the kind of misperceptions of public information necessary to reconcile models of forecaster behavior with the observed overresponse to public information. The next section then proposes a particular model, that of naïve overconfidence, which falls within this class, and is consistent with all of the stylized facts.

Beliefs about Public Information: Consider once more individual i's forecast in (4.2)

$$f_i = (1-w)\,\mathbb{E}[\theta \mid \mu_i, y] + w x_i,$$

where the the weight accorded to private information does not necessarily equal its mean-squared optimal counterpart, $w \neq w_\star$. But suppose now that, in addition to any potential misuse of private information (relative to mean-squared optimal), the forecaster simultaneous misinterprets the public signal y. Specifically, suppose that he perceives *the realized public signal* y to equal *the perceived public signal* y_p. In that case, we can write this individual's forecast as

$$f_i = (1-w)\,\mathbb{E}[\theta \mid \mu_i, y_p] + w x_i, \tag{5.1}$$

where, with a slight abuse of notation, $\mathbb{E}[\theta \mid \mu_i, y_p]$ denotes the conditional expectation of θ based on the realized public signal y being treated as the perceived signal y_p. We also for simplicity assume that μ_i is consistent $\mu \mid \theta \sim \mathcal{N}\left(\theta, \tau_\mu^{-1}\right)$ and independent of all other disturbances. It then follows that this forecaster's forecast error will be

$$\theta - f_i = \theta - w x_i - (1-w)\,\mathbb{E}[\theta \mid \mu_i, y_p]. \tag{5.2}$$

Taking conditional expectations based upon *realized public information* then shows that

$$\begin{aligned}\mathbb{E}\left[\theta - f_i \mid y\right] &= \mathbb{E}\left\{\theta - \mathbb{E}\left[\theta \mid \mu_i, y_p\right] \mid y\right\} \\ &= \mathbb{E}\left\{\mathbb{E}\left[\theta \mid \mu_i, y\right] - \mathbb{E}\left[\theta \mid \mu_i, y_p\right] \mid y\right\} \neq 0. \end{aligned} \tag{5.3}$$

Unlike in (4.4) above, the Law of Iterated Expectations no longer implies orthogonality between individual forecast errors and public information. This is because $\mathbb{E}\left[\mathbb{E}\left[\theta \mid \mu_i, y_p\right] \mid y\right] \neq$

$\mathbb{E}[\theta \mid y]$. A relationship therefore arises between individual forecast errors and consensus realizations, $\delta \neq 0$. The misperceptions of public information breaks the implication of the Law of Iterated Expectation that forecast errors are orthogonal to public information.[16]

Overresponses to Public Information: Generically, there are two means by which forecasters can overrespond to observed public information. They can (a) overstate the precision of the information that they receive and (b) over-interpret moves in the fundamental from said information. To appreciate the conditions under which these two sources of misinterpretation of public news can result in overresponses, suppose *the realized public signal* equals

$$y = \kappa\theta + \xi, \tag{5.4}$$

where $\kappa \in \mathbb{R}_+$, while *the perceived public signal* is

$$y_p = \theta + \xi. \tag{5.5}$$

An individual's rational forecast based on μ_i and y then becomes

$$\mathbb{E}[\theta \mid \mu_i, y] = (1 - k_\star)\mu_i + k_\star y, \quad k_\star = \frac{\kappa^2\tau_\xi}{\tau_\mu + \kappa^2\tau_\xi} \times \frac{1}{\kappa} \tag{5.6}$$

while the corresponding forecast based on y_p is

$$\mathbb{E}[\theta \mid \mu_i, y_p] = (1 - k)\mu_i + ky, , \quad k = \frac{\tau_\xi}{\tau_\mu + \tau_\xi} \times 1. \tag{5.7}$$

Thus, in this case, (5.8) simply becomes

$$\mathbb{E}[\theta - f_i \mid y] = (k_\star - k)\,\frac{\kappa - 1}{\kappa} y \neq 0. \tag{5.8}$$

As argued above, there are two important distinctions between the public information that forecasters receive and what they believe the receive, and hence between (5.6) and (5.7).

First, the realized public signal y has a different precision that the suspected public signal y_p. Its precision about θ is $\mathbb{V}[\theta \mid y] = (\kappa^2\tau_\xi)^{-1}$ rather than $\mathbb{V}[\theta \mid y_p] = \tau_\xi^{-1}$, as suspected. This, all else equal, cases the weight on y in (5.7) to differ from its rational counterpart. Specifically, when forecasters are overconfident in the precision of the public information, $\kappa > 1$, they all else equal overrespond to the public signal. This can be seen from how $k_\star$, all else equal,

[16] Footnote about data being seen both by econometrican and forecasters.

increases above k when one increases. As a result, a negative correlation between individual forecast errors and the realized public signal arises, $\delta < 0$.

Second, because $\kappa \neq 1$, the realized public signal, however, also has a different loading on the fundamental than the suspected public signal. A move of $d\theta$ cause a movement of $\kappa d\theta$ in y relative to $d\theta$ in y_p. This, in addition, causes the weight on the realized public signal y in (5.7) to differ from its rational counterpart. In particular, when forecasters ‘ underestimate the extent to which movements the public signal reflect the fundamental ($\kappa > 1$), they deduce excessive movements in the fundamental. A move of $d\theta$ causes the perception of movement equal to $\kappa d\theta > d\theta$. This, in turn, causes forecasters to overrespond to the public signal, and hence for negative correlation to arise between forecasts errors and public information, $\delta < 0$.

Focusing on the relevant case where the public signal is informative relative to individual priors ($\tau_\xi > \tau_\mu$), and on small departures from the noisy rational expectations benchmark ($\kappa > 1$ but close to 1), misperception of public signals thus leads to overreaction whenever forecasters perceive consensus forecasts either to react less to fundamentals than they actually do (\kappa_{i}<1), or to be more precise than they actually are (\sigma_{i}<1), or both. In the first case, when the public signal reacts more to fundamentals than forecasters think, they infer too much fundamental news from any given observed movement in consensus. In the second case, when consensus noise is larger than forecasters think, they give too much weight to consensus forecasts.

At its heart, whether forecasters' over- or underrespond to public information depends on the weight accorded to public news in individuals' forecast. This weight, in turn, is a function of two components: forecasters perceived accuracy of public information and how forecasters interpret the public news. On balance, forecasters that misperceive the information that they receive commit both types of mistakes. But to create the observed overresponse to public information that we documented in Section 3 forecasters either need to overestimate the precision of their information or under-appreciated its responsiveness to the forecasted variable. In the next section, we analyze a model of forecaster behavior that results in precisely this type of misperception.

6 Naïve Overconfidence

In this section, we show how a simple model of *naïve overconfidence* can reconcile our empirical results. We call "naïve" those individuals that are not only overconfident in their own information but also wrongly think that their information is superior to that of others. We therefore merge the two related but distinct notions of overconfidence commonly used in the psychology literature (Moore and Healy, 2008). We refer to the first type of overconfidence as *overprecision* and the latter as *overplacement* (Benoît *et al.*, 2015). We then show how such naïvety, combined with the realization that most public information reflects endogenous choices of others, can reconcile our stylized facts.

6.1 A Model of Naïve Overconfidence

Consider a two-period version of the baseline model from Section 2. We will only require an explicit account of time because we later will equate the public information that forecasters observe with last period's consensus estimate. At the start of each period each forecaster $i \in [0, 1]$ receives the private signal x_{it} about the random variable $\theta \sim \mathcal{N}(\mu, \tau_\theta)$

$$x_{it} = \theta + \epsilon_{it}, \quad t = \{1, 2\},$$

where $\epsilon_{it} \sim \mathcal{WN}\left(0, \tau_x^{-1}\right)$ is independent of θ and $\mathbb{E}[\epsilon_{it}\epsilon_{jh}] = 0$ for all $j \neq i$ and all $t, h = \{1, 2\}$. Each forecaster i is naïvely overconfident in the precision of his own information. The forecasters believes the precision of his private information to be $\tau_x' > \tau_x$, whereas he believes the precision of other forecasters private information simply to be τ_x.[17] Each forecaster thus, in addition to being overconfident, wrongly thinks that his private information is better than others. This differentiates the type overconfidence studied here from that explored in Section 4.[18] Last, at the start of the second-period, each forecaster also observes a noisy signal of last period's consensus estimate[19]

$$f = \int_0^1 f_{i1} + \xi, \tag{6.1}$$

where f_{i1} denotes i's first-period forecast and $\xi \sim \mathcal{N}\left(0, \tau_y^{-1}\right)$ is independent of all other shocks.

Individual Forecasts: Consider forecaster i's first-period forecast

$$f_{i1} = vx_{i1} + (1 - v)\,\mu, \quad v = \frac{\tau_x'}{\tau_x' + \tau_\theta}, \tag{6.2}$$

where v exceeds the mean-squared error optimal weight on private information, $v_\star = \frac{\tau_x}{\tau_x + \tau_\theta}$ because of the forecasters' overconfidence in the precision of his own information.

Let $\mu_i = f_{i1}$ now denote forecaster i's prior expectation at the start of the second-period. To derive an individual's $t = 2$ forecast, we first need to differentiate between two different

[17] Equating forecasters beliefs about the precision of other's private information to its actual precision simplifies the below derivations. But what is central for the below argument is simply that forecasters believe themselves to be superior and hence distort their interpretation of public information.

[18] Moore and Healy (2008) review the literature that document both overprecision and overplacement across a wide variety of experimental tasks.

[19] The reason for the introduction of the shock in (6.1) is purely technical: the important role it plays is to limit forecasters' ability to infer the true value of θ from the observation of f. The use of "non-invertibility" shocks has been standard in the noisy rational expectations literature since Hellwig (1980).

consensus estimates: (i) the *realized* consensus forecast f, and (ii) the *suspected* consensus forecast f_s. The former measures the *actual* noisy signal of consensus in (6.1)

$$f = \int_0^1 f_{i1} + \xi = (1 - v)\,\mu + v\theta + \xi. \tag{6.3}$$

The latter, by contrast, measures the consensus forecast that forecasters *believe* they observe

$$f_s = (1 - v_\star)\,\mu + v_\star\theta + \xi. \tag{6.4}$$

f and f_s differ due to the failure of forecasters to internalize the overconfidence of others; that is, that all forecasters attach a weight of $v > v_\star$ to their own private information at $t = 1$.

Forecasters who observe a realized consensus forecast of f in (6.3) thus treat it *as-if* it was generated by f_s in (6.4). As a result, the public signal that forecasters infer from consensus becomes

$$y = \frac{1}{v}\left[f - (1 - v)\,\mu\right] = \frac{v}{v_\star}\theta + \frac{v_\star - v}{v_\star}\mu + \frac{1}{v_\star}\xi. \tag{6.5}$$

Nevertheless, the forecasters believe that they observe

$$y_s = \theta + \frac{1}{v_\star}\xi. \tag{6.6}$$

Conditional on the common prior, y in (6.5) and y_p in (6.6) fall within the class studied in (5.4) and (5.5) with $\kappa = \frac{v}{v_\star}$. As a result, there are two important distinctions between y and y_s, both of which are driven by forecasters inherent misperception about others caused by overplacement.

First, the realized consensus outcome is more precise than the the suspected one. Conditional on μ, the precision of y about θ is $v^2\tau_\xi$, while the precision of of y_s is only $v_\star^2\tau_\xi$, where $v > v_\star$. Since all forecasters respond more to their private information than expected, consensus embeds more of the truly new (private) information that forecasters can learn from each other. This increases the informativeness of the consensus estimate.

Second, and related, naïve forecasters deduce excessive values of the fundamental from consensus. The realized consensus loads onto the fundamental with $v/v_\star > 1$, while the suspected consensus only loads onto the fundamental with one. A movement of $d\theta$ in the fundamental, thus, causes forecasters to all else equal believe in a moment equal to $(v/v_\star)\,d\theta$, based on the observation of consensus alone.

Put succinctly, the misperceptions about others inherent to our notion of naive overconfidence both causes forecasters to underestimate the precision of consensus as well as to wrongly infer information from it.

We are now ready to state forecasters' second-period forecast:

$$f_{i2} = (1-w)\,\mathbb{E}\left[\theta \mid y_s, \mu_i\right] + w x_{it}, \tag{6.7}$$

where $\mathbb{E}\left[\theta \mid y_s, \mu_i\right]$ denotes the conditional expectation of θ based on the realized public signal y being treated as if it were y_s, in addition to the individual-specific prior μ_i. We note that $w = \frac{\tau_x'}{\tau_\theta + 2\tau_x' + v'^2\tau_\xi} \in (w_\star, 1)$ when $\tau_\theta < \tau_x$.[20]

Misperceptions About Others: The forecasts in (6.7) are consistent with our empirical results from Section 3. First, because of the presence of noisy private information, on average forecasters underrespond to new information, as measured by the average forecast revision. Second, despite this underresponsiveness at the average level, at the individual level, forecasters can instead overrespond because of their overconfidence in the precision of their private information. And last, due their misperception about others, a negative relationship can arise between the realized consensus estimate and individual forecast errors.

Proposition 3. *Consider individual $i \in [0,1]$'s second-period forecast*

$$f_{i2} = (1-w)\,\mathbb{E}[\theta \mid \mu_i, y_s] + w x_i, \tag{6.8}$$

Then, if $\tau_\theta < \min\left\{\tau_x, \frac{\tau_x'}{\tau_x} v_\star^2 \tau_\xi\right\}$, $b > 0$ in (2.4), $\beta < 0$ in (2.6) and $\delta < 0$ in (2.7).

The first and second result in Proposition 3 follow from Proposition 2 and that $w \in (w_\star, 1)$ and therefore omitted here. The final result is a direct consequence of (6.5) and (6.6) falling within the class studied in Section 5 with $\kappa = \frac{v}{v_\star}$.

Combined, the misperceptions about others at the heart of naïve overconfidence and the presence of endogenous public information breaks the implication of the Law of Iterated Expectation that forecast errors are orthogonal to public information. In fact, when $\tau_\theta < \min\left\{\tau_x, \frac{\tau_x'}{\tau_x} v_\star^2 \tau_\xi\right\}$

$$\begin{aligned}
\mathbb{E}\left[\theta - f_i \mid y\right] &= (1-w)\,\mathbb{E}\Big[\mathbb{E}\left[\theta \mid \mu_i, y\right] - \mathbb{E}\left[\theta \mid \mu_i, y_s\right] \mid y\Big] \\
&= \frac{(1-w)(1-\frac{v}{v_\star})v_\star^2\tau_\xi}{\left(\frac{v}{v_\star}\right)^2 v_\star^2\tau_\xi + \tau_\mu}\left[\frac{v}{v_\star} + \left(1 + \frac{v}{v_\star}\right)\frac{\tau_\mu}{\tau_\mu + v_\star^2\tau_\xi}\,\frac{\frac{v}{v_\star}(1-\frac{v}{v_\star})v_\star^2\tau_\xi - \tau_\theta}{\left(\frac{v}{v_\star}\right)^2 v_\star^2\tau_\xi + \tau_\theta}\right] y < 0
\end{aligned}$$

since $v/v_\star > 1$ and the term inside the bracket is positive. As a consequence, an overresponse to

[20] The necessary and sufficient condition for $w \in (w_\star, 1)$ in (6.7) is $\tau_\xi\tau_x^2(\tau_\theta - \tau_x) < \tau_\theta^4 + 3\tau_\theta^3\tau_x + 3\tau_\theta^2\tau_x^2 + \tau_\theta\tau_x^3 + \tau_\xi\tau_x^3 - \tau_\xi\tau_\theta\tau_x^2$. The reason for this dependence on the precision of consensus and the prior is simple: In order to have, for instance, underresponses at the average level, we need forecasters to attach sufficient emphasis onto their own private information, to counteract the overresponses that will occur to consensus, for example. This, in turn, requires the consensus and the prior not to be excessively informative.

public information arises in conjunction with an under- and overresponse of forecast revisions at the average and individual level, respectively.

Overresponses to Public Information: As argued above, there are two important distinctions between the public information that forecasters receive and what they believe the receive.

First, the realized public signal y is more precise than the suspected public signal y_s. Its precision about θ is $v^2\tau_\xi$ rather than $v_\star^2\tau_\xi$ as suspected, where $v > v_\star$. This, all else equal, causes the weight on y in (6.8) to be too small. When forecasters are overconfident in the precision of their own information relative to others, they underestimate the extent to which consensus embeds the truly new (private) information that forecasters can learn from one another. This, in turn, causes forecasters to underrespond to consensus realizations, and as a result a positive correlation between individual forecast errors and consensus realizations arises, $\delta > 0$.

Second, because $v_\star/v < 1$, the realized public signal, however, also has a larger loading on the fundamental than the suspected public signal. A move of $d\theta$ cause a movement of $v/v_\star d\theta$ in y relative to $d\theta$ in y_s. This, by contrast, causes the weight on the realized public signal y in (4.5) to be too large, and hence for $\delta < 0$, all else equal. When forecasters are overconfident in their own information vis-à-vis others, they underestimate the extent to which movements in consensus correctly reflect the fundamental. This, in turn, causes forecasters to deduce excessive values of the fundamental from consensus realizations, and hence for forecasters to overrespond to consensus. As a result, $\delta < 0$ when consensus is sufficiently informative, $\tau_\theta < \frac{\tau_x'}{\tau_x} v_\star^2 \tau_\xi$.

At its heart, whether forecasters over- or underrespond to public information depends entirely on the weight accorded to public information in individuals' forecast. This weight, in turn, is a function of two components: forecasters perceived accuracy of the public information and how forecasters interpret the public news. Naively overconfidence causes forecasters to mistake both. In the next section, we analyze the forces that determine the relative strength of these two erroneous inferences, and hence whether forecasters over- or underrespond to public information.

6.2 Comparative Statics

The previous section showed how a model of naïve overconfidence is qualitatively consistent with all three stylized facts about forecaster behavior documented in Section 3. In this subsection, we explore the capacity of the model to match the stylized facts for different parameter values. This will create insight about the precise forces that allow our model to reconcile our empirical results.

To do so, we use a simulated method of moments procedure to choose parameter values that minimizes a weighted average of differences between a vector of moments implied by the model and those found in the data. Apart from the regression coefficients β and δ, we also include the standard deviation of forecast errors across all forecasters and time periods among our target moments, presumably one of the most important statistics for economic forecasts.[21] For the estimation, the criterion we choose to minimize is

$$\Lambda(\tau) \quad = \quad [\hat{m} - m(\tau)]' W^{-1} [\hat{m} - m(\tau)],$$

where $\hat{m}$ is a vector of target moments of the SPF data and $m(\tau)$ is the vector of simulated moments as a function of the parameter vector $\tau = (\tau_{x1}, \tau_{x2}, \tau_{\xi}, \tau_{\theta})$ and, when applicable, the overconfidence parameters. For the estimation we normalise the signal precision τ_θ to 1, and assume that both actual and perceived precision of forecasters' signals are constant over time, equal to τ_x, τ_x' respectively in both periods. For our estimation we use a diagonally weighted minimum distance procedure, corresponding to a weighting matrix W that has the variances of the moments on the diagonal and is zero everywhere else.

Table II: SMM Estimation: Inflation Forecasts

	β	δ	σ_{error}	$\sqrt{\tau_x}$	$\sqrt{\tau_\xi}$	$\sqrt{\tau_\theta}$	$\sqrt{\tau_{x'}}$
Data	-0.19	-0.10	0.98				
Naive Overc.	-0.15	-0.09	0.98	0.50	5.00	1.00	1.00

Table II presents the results, [where for ease of interpretation we report standard deviation $\sigma = \frac{1}{\sqrt{\tau}}$ of random variables and priors instead of precision τ]. The model of forecaster superiority captures all three data moments well. We estimate private signals to be rather noisy ($\tau_x = 0.25$) and consensus-noise to be small ($\tau_\xi = 25$). At a level of overconfidence that halves the perceived standard deviation of private signals, the model predicts almost perfectly both the volatility of forecast errors and the overresponse to public information as captured by the coefficient δ, but somewhat underpredicts β, the coefficient summarizing the overresponse to private information. Appendix D presents the corresponding estimates for US SPF forcasts for real GDP growth and CPI inflation. The model cannot capture the large overresponse that we estimate for GDP growth forecasts, but otherwise captures the estimated coefficients

[21] We calculate the standard deviation using a simple resampling bootstrap procedure with $N = 100,000$ resamples.

well, with similar estimated levels of overconfidence, small estimated noise in consensus and somewhat more precise private signals.

Figure 2 illustrates how the predicted regression coefficients β and δ vary as model parameters change from their estimated levels in Table II. For this, we vary individual parameters along the bottom axis of the three panels, but keep the relative overconfidence $\frac{\sqrt{\tau_x'}}{\sqrt{\tau_x}}$ and all other parameters at their estimated values. The top panel shows that when private signals are more precise than our estimate, the regression coefficient β, summarizing the overresponse to private information, can actually turn positive. This is because β captures the overreaction to private information in period two at the same time as the correction in the first-period overconfident forecast. This correction becomes more important the more precise is the additional information received in the second period. [...] As the top panel also shows, when private signals are more noisy than our estimate, and τ_x thus lower, the regression coefficient δ, summarizing the response to public information, turns positive. This is because, as Section 6 shows, δ captures the effect of two simultaneous departures from the rational forecast: a wrong interpretation of the fundamental θ (to which consensus responds more than forecasters think), and a wrong perception of its precision (which is higher than forecasters think). The former effect by itself implies an overreaction to consensus ($\delta < 0$), while the latter implies an underreaction ($\delta > 0$), relative to a rational forecast with $\delta = 0$. Which of the two effects dominates depends crucially on signal noise and other parameters. [...]

The middle panel of Figure 2 shows that positive values of δ also arise when consensus noise increases (τ_ξ falls) from its estimated value. This is because the misperception of consensus precision is increasing in the level of consensus noise.

Finally, the bottom panel of Figure 2 shows a U-shaped relationship of both coefficients as overconfidence, and thus the perceived precision of private signals τ_x' relative to the true precision τ_x, rise. This is, on the one hand, because whenever private signals are perfectly revealing ($\tau_x \longrightarrow \infty$) there is no room for overresponses to either public or private information as both rational and superior forecasts equal the average private signal. As private signals loose their information, on the the other hand, both rational and superior forecasts equal the weighted average of prior and consensus, again eliminating any room for overresponse.

Figure 2: Coefficient Estimates

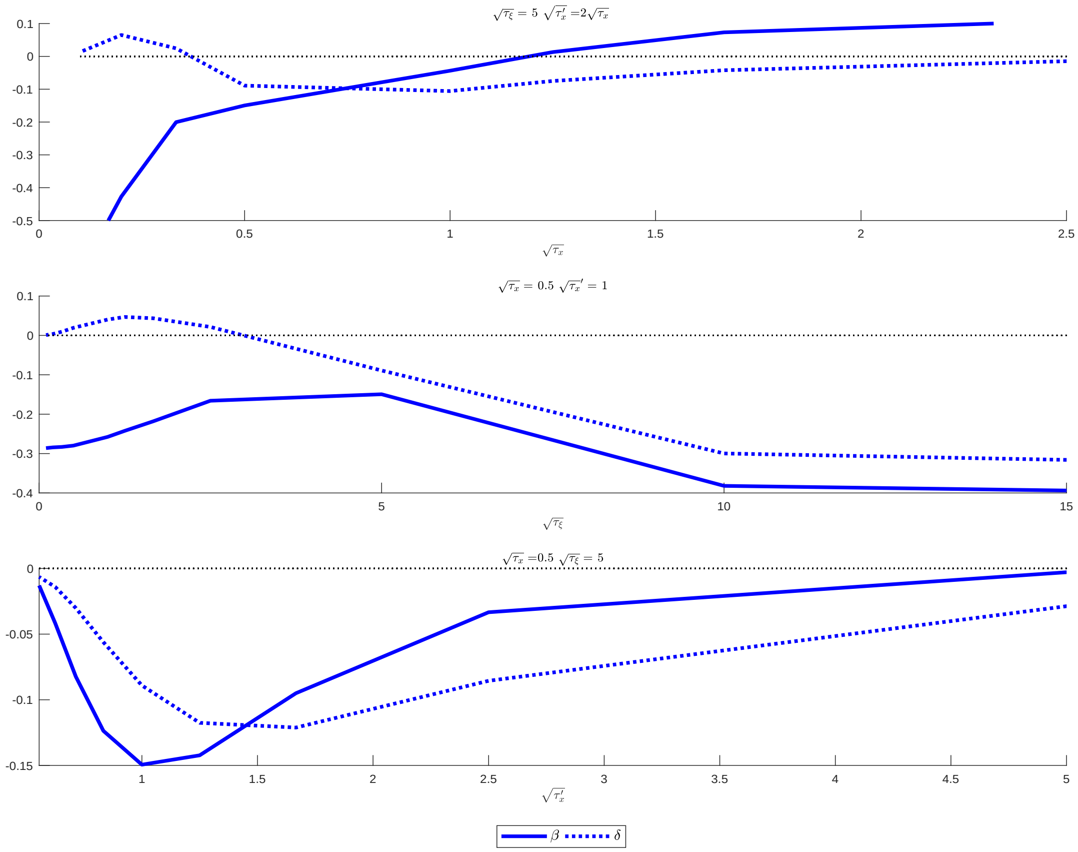

The figure depicts the regression coefficients β_1 and δ_1 to δ_3 in equation (**??**) estimated on simulated data from the LIRE benchmark of Section **??** ("Rational"), and the two simple alternative models of strategic diversification ("Strategic") and overconfidence ("Overconfident") presented in Section **??**. In the simulations τ_θ is normalized to 1, $\sqrt{\tau_\xi}$ equals $\frac{1}{2}$ and the overconfidence parameter τ'_{xt} is set to $4\tau_x$.

7 Conclusion

A considerable debate has arisen about the best model of expectation formation. Previous research has shown that average forecasts across a variety of surveys are consistent with models of dispersed, imperfect information and rational information use (see, for example, Coibion and Gorodnichenko, 2015). In this paper, we in contrast have explored the implications of such models for the behavior of individual forecasts conducted by professional forecasters.

We have documented how the behavior of individual forecasts of inflation contradict simple versions of noisy rational expectations. Specifically, individual forecasts overrespond to new information on average, and in particular to readily available public signals, such as consensus estimates. We have shown how such overresponses violate a basic tenet of noisy rational expectations, the law of iterated expectation. As a result, such overresponses create correlation between individual forecast errors, on the hand, and individual forecast revisions and past public information, on the other hand. Importantly, we have illustrated how such overresponses also extend to other variables than inflation and across different countries.

Several standard theories of forecaster behaviour are consistent with the average overresponses to new information, in addition to average forecasts that resemble those that arise from noisy rational expectations. But as we have documented above, such models are also often inconsistent with the simultaneous overresponse to readily available public signals, such as consensus estimates. In place, we above proposed a new model of naïve overconfident forecasters consistent with the stylized facts.

We have called naïve those forecasters that believe that their information is not only better than it truthfully is but also better than that available to their their competitors. These biases of, respectively, "overprecision" and "overplacement", have been extensively documented in both experimental and other contexts (Moore and Healy, 2008). Combined, they entail that forecasters interpret a given change in public information as an average of weak responses to news by poorly informed forecasters, while in fact it results from forceful responses by overconfident forecasters to a small movement in fundamentals. We showed how this model can rationalize the observed forecast data both qualitatively and quantitatively.

More generally, our theory has highlighted the possible sources of suboptimal responses to public information. Mathematically, these can arise whenever there is a difference between perceived and true expectations conditional on public information. Such differences could, in principle, result from three sources. First, from differences in the public information available to people relative to researchers. Second, from departures from Bayes' Rule in processing information. And last, third, whenever public signals are endogenous, from agents' misperception of the information content of public information caused by because they misunderstanding the behavioral rules of other agents and the equilibrium mechanisms that aggregate them into public signals.

Our empirical estimates have been about the response of professionals to readily available, salient and relevant information, consisting of simple averages of individual forecasts that are not subject to revisions. This suggests that, in this particular context, differences in information sets, failures to update expectations as suggested by Bayes' Rule, and mi-perceptions of equilibrium aggregation mechanisms are unlikely to explain the observed correlations. This is why we built a theory based on behavioral biases that preserve the use of Bayes' Rule in the update of information, but imply misperceptions about the informativeness of public information. In other contexts, however, other sources of misperception may be relevant

Last, our results entail specific opportunities for further research. First, it seems worthwhile to explore if our empirical results also translate to other agents than professional forecasters, such as firms. In contemporaneous work, Bordalo *et al.* (2018a) indeed document overresponses to new information on average across a wide variety of different agent. But it would seem central for our theory of naïve overconfidence to also document whether such overresponses extend to specific pieces of endogenous public information. Second, Kohlhas and Walther (2018) show how apparent overresponses to public information can arise rationally, as the outcome of forecasters' optimal information choice under dispersed, noisy information. Attempting to disentangle the extent to which measured overresponses to public information arise from behavioural overconfidence or optimal information choice seems a natural next step.

References

Anderson, B. D. and Moore, J. B. (1979). Optimal filtering. *Englewood Cliffs*, **21**, 22–95.

Andrade, P. and Le Bihan, H. (2013). Inattentive professional forecasters. *Journal of Monetary Economics*, **60** (8), 967–982.

Benoît, J.-P., Dubra, J. and Moore, D. A. (2015). Does the better-than-average effect show that people are overconfident?: Two experiments. *Journal of the European Economic Association*, **13** (2), 293–329.

Bhattacharya, S. and Pfleiderer, P. (1985). Delegated portfolio management. *Journal of Economic Theory*, **36** (1), 1–25.

Bordalo, P., Gennaioli, N., Ma, Y. and Shleifer, A. (2018a). Overreaction in macroeconomic expectations.

—, — and Shleifer, A. (2018b). Diagnostic expectations and credit cycles. *The Journal of Finance*, **73** (1), 199–227.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, **105** (8), 2644–78.

Croushore, D. and Stark, T. (2001). A real-time data set for macroeconomists. *Journal of econometrics*, **105** (1), 111–130.

Croushore, D. D. (1993). Introducing: the survey of professional forecasters. *Business Review-Federal Reserve Bank of Philadelphia*, **6**, 3.

— (1997). The livingston survey: still useful after all these years. *Business Review-Federal Reserve Bank of Philadelphia*, **2**, 1.

Cutler, D. M., Poterba, J. M. and Summers, L. H. (1990). Speculative dynamics and the role of feedback traders.

Daniel, K., Hirshleifer, D. and Subrahmanyam, A. (1998). Investor psychology and security market under-and overreactions. *the Journal of Finance*, **53** (6), 1839–1885.

Dovern, J., Fritsche, U., Loungani, P. and Tamirisa, N. (2015). Information rigidities: Comparing average and individual forecasts for a large international panel. *International Journal of Forecasting*, **31** (1), 144–154.

Ehrbeck, T. and Waldmann, R. (1996). Why are professional forecasters biased? agency versus behavioral explanations. *The Quarterly Journal of Economics*, **111** (1), 21–40.

Ericson, W. A. (1969). A note on the posterior mean of a population mean. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 332–334.

Fuhrer, J. C. (2015). Expectations as a source of macroeconomic persistence: an exploration of firms' and households' expectation formation.

Garcia, J. A. (2003). An introduction to the ecb's survey of professional forecasters.

Graham, J. R. (1999). Herding among investment newsletters: Theory and evidence. *The Journal of Finance*, **54** (1), 237–268.

Granger, C. W. (1969). Prediction with a generalized cost of error function. *Journal of the Operational Research Society*, **20** (2), 199–207.

Hellwig, M. F. (1980). On the aggregation of information in competitive markets. *Journal of economic theory*, **22** (3), 477–498.

Kahneman, D. and Tversky, A. (1973). On the psychology of prediction. *Psychological review*, **80** (4), 237.

Kohlhas, A. N. and Walther, A. (2018). Asymemtric attention.

Lamont, O. A. (2002). Macroeconomic forecasts and microeconomic forecasters. *Journal of economic behavior & organization*, **48** (3), 265–280.

Laster, D., Bennett, P. and Geoum, I. S. (1999). Rational bias in macroeconomic forecasts. *The Quarterly Journal of Economics*, **114** (1), 293–318.

Marinovic, I., Ottaviani, M. and Sorensen, P. (2013). Forecasters objectives and strategies. *Handbook of economic forecasting*, **2**, 690–720.

Moore, D. A. and Healy, P. J. (2008). The trouble with overconfidence. *Psychological review*, **115** (2), 502.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, pp. 315–335.

Ottaviani, M. and Sørensen, P. N. (2006). The strategy of professional forecasting. *Journal of Financial Economics*, **81** (2), 441–466.

Pesaran, M. H. and Weale, M. (2006). Survey expectations. *Handbook of economic forecasting*, **1**, 715–776.

Scharfstein, D. S. and Stein, J. C. (1990). Herd behavior and investment. *The American Economic Review*, pp. 465–479.

Welch, I. (2000). Herding among security analysts. *Journal of Financial economics*, **58** (3), 369–396.

Zarnowitz, V. (1985). Rational expectations and macroeconomic forecasts. *Journal of Business & Economic Statistics*, **3** (4), 293–311.

— and Lambros, L. A. (1987). Consensus and uncertainty in economic prediction. *Journal of Political economy*, **95** (3), 591–621.