Strategic Forecasting

Francesco Amodeo*

March 2025

Abstract

This study proposes a novel approach to interpreting conventional tests of the Full Information Rational Expectations (FIRE) hypothesis, based on the distinctions between forecasts and expectations. First, I argue that these two objects coincide under restrictive conditions that are unrealistic in the context of professional forecasting. Evidence from reduced-form analysis suggests that forecasters pursue strategic incentives when responding to surveys, contaminating their responses and making their use as measures of expectations misleading. Second, I leverage this distinction to introduce a new parsimonious model of forecast formation that is related to rational inattention and sparsity-based models of bounded rationality, but that is exempt from their complications, including the dependence on Gaussian information. The proposed framework employs a global game structure featuring public and private information, as well as the strategic behavior of forecasters, and demonstrates that strategy explains the anomalies in Coibion and Gorodnichenko (2015)-type regressions. Finally, I exploit the model's transparency to develop and apply a formal testing methodology that validates the strategic channel, advising caution in using surveys to elicit expectations.

Keywords: Expectations, Forecasts, Information, Strategy, Surveys
JEL Codes: C53, D83, D84, E13, E17, E31, E47

*UC San Diego, famodeo@ucsd.edu
I am grateful to James Hamilton for his support and guidance. I have benefited from conversations with Giacomo Rondina, Allan Timmermann, Johannes Wieland, Juan Herreño, Nir Jaimovich, Yuriy Gorodnichenko, Raffaella Giacomini, Luca Gemmi, Julian Kozlowski, Ali Uppal, Snehal Banerjee and seminar participants at UC San Diego, the Rady School of Management, the UCSD School of Global Policy and Strategy, the 2022 NBER *Behavioral Macroeconomics*, the 2024 *Sailing the Macro* Workshop, the 2025 Midwest Macroeconomics Meeting and the EAYE 2025 Annual Meeting. This paper previously circulated under the title *"Expectations and Forecasts: The Role of Strategic Incentives"* (first draft: September 2022).

1 Introduction

Expectations play a central role in macroeconomics and finance. However, the mechanics underlying their formation process remain elusive. A major barrier to the understanding of their dynamics is the unobservability of agents' expectations, which poses a challenge to testing different models of belief formation. A ground-breaking shift in the discourse has been the incorporation of survey forecast data, offering a micro-level perspective in the analysis of expectations. However, robust counterintuitive patterns have emerged in the past decade, indicating the coexistence of under and overreaction to news respectively at the aggregate and at the individual level.

Coibion and Gorodnichenko (2012),(2015) first showed the empirical regularities that separate the survey evidence from the ideal paradigm of Full Information Rational Expectations (FIRE), exploring the correlation between ex post forecast errors and ex ante forecast revisions at the *consensus* level. They proposed a theory of information rigidities based on the benchmark models of Mankiw and Reis (2002) and Woodford (2001), attributing the deviations from the frictionless scenario to *noisy* or *sticky* information, while maintaining the assumption of individual rationality. In this fashion, a whole branch of the literature has corroborated and expanded these results by examining different features of information dissemination (Andrade and Le Bihan (2013), Sarte (2014)), acquisition (Carvalho and Nechio (2014), Kohlhas and Walther (2021)) and processing (Cavallo et al. (2017), Angeletos et al. (2021), Farmer et al. (2021)).

Starting with Bordalo et al. (2020), a concurrent set of studies has identified a puzzling predictability in *individual* forecast errors, casting fundamental doubts on the rationality of forecasters. Papers in this literature motivate the empirical anomalies with psychologically founded models that incorporate a plethora of behavioral biases and cognitive failures, such as salience and diagnosticity (Bordalo et al. (2018)), over-extrapolation (Barberis et al. (2018), Fuster et al. (2010), Fuster et al. (2012)), over-confidence (Scheinkman and Xiong (2003), Broer and Kohlhas (2022)) and the effect of subjective experience (Malmendier and Nagel (2016), Da Silveira et al. (2020)).

This paper advances a parsimonious, testable and transparent framework, based on the distinction between expectations and forecasts, that reconciles the apparent anomalies without resorting to theories of cognitive failure or information contamination. By neglecting strategic behavior, previous studies have inadvertently built these motives into frameworks in which strategy is absent, therefore biasing the existing tests used to reject rationality and full information. Intuitively, incentives not aligned with the maximization of accuracy – as in the case of strategic competition among forecasters – are not captured by forecast errors, the ubiquitous dependent variable in FIRE tests. Thus, inference drawn from such regressions will inevitably be biased due to the omission of a dimension of optimization in the forecaster's problem.

Adopting a new approach to modeling information, this study departs from the typical formulation where agents observe a signal about a state variable that is contaminated by random noise, with no relation of the signal to the observable fundamentals of the economy. Instead, it represents private information as an actual component of the realization, one that is accurately anticipated only by some individuals, in a related fashion to Chahrour and Jurado (2018). In this framework, an imprecise (*noisy*) signal corresponds to the foreknowledge of only a minor component of the actual realization. This representation of the nature of private information allows the researcher to describe the forecasting problem organically to the economic environment.

Integrating strategy in forecasters' objectives explains the consensus underreaction, the individual overreaction, as well as some new summaries of the data I present. The model allows making predictions for these and potentially any other empirical moment. Importantly, the transparent nature of the framework allows to test formally the predictions of the information structure itself, confirming the statistical robustness of the conjectured mechanisms. Adapting the methodology from Hansen (1982), I show that in the overwhelming majority of scenarios, across variables and horizons, the data fails to reject model-specific restrictions, validating the necessity of strategic considerations in explaining the observed patterns.

The introduction of strategic incentives into the macro-forecasting discourse is a novel development. To the best of my knowledge, this is among the very first papers to study in depth strategic motives within the domain of survey evidence about expectations formation processes. Previous studies suggest strategy affects survey response, yet the direction of this effect remains context-dependent. In some settings, reputational concerns lead forecasters to herd on the consensus (Gallo et al. (2002), Lamont (2002), Hong and Kacperczyk (2010)), while in others they may deliberately deviate to gain attention or rewards (Laster et al. (1999), Ottaviani and Sørensen (2006a), Bernhardt et al. (2006)). Introducing additional moments in the characterization of the data generating process of the Survey of Professional Forecasters, my results favor the latter interpretation in the context of macroeconomic forecasting, revealing differentiation forces able to offer a unified rationale for the apparent anomalies.

In important contemporaneous work, Gemmi and Valchev (2023) contribute with insightful considerations into the strategic aspects of forecasting, while operating within the standard framework of *noisy information* models. In contrast, this paper advances a new framework based on an explicit characterization of the process governing the forecasted variable. Rather than treating private information as an exogenous signal contaminated by random noise, this approach conceptualizes it as a component of the future realization that some individuals know in advance. This approach presents three main advantages: (i) it allows for formal statistical testing of the restrictions imposed on the model, key to assessing the ability of the model to explain the data, a feature not present in the overwhelming majority of studies in the literature; (ii) it achieves a neat separation between the information and the strategic spheres, favoring the interpretability and sharpness of the economic mechanisms; (*iii*) it *does not* depend on assumptions like Gaussian signals, assumption that I show to be empirically implausible and yet key for tractability and estimation in virtually all other models. Moreover, Gemmi and Valchev (2023) leverage the presence of public information, while in this study, although public signals can be present, the key insights hold irrespectively of the information environment, and the main forces operate analogously with purely idiosyncratic information. The emergence of similar insights across the two different approaches reinforces the conclusions of each.

This novel framework, affine in spirit to the affiliated signals used in the auction literature (Milgrom and Weber (1982)), predicts that when presented with private information, agents wanting to *stand out* from the crowd will rationally overweigh it to differentiate themselves. In the preferred baseline specification, the strength of the strategic motive estimated using the data from the Federal Reserve's Survey of Professional Forecasters is large and significant: across variables and horizons, the median estimates of strategic reasoning suggest that it accounts for more than half of the importance of accuracy in respondents' objective function.

Moreover, additional implications of the strategic model help in distinguishing it from both rational and irrational frameworks with similar implications. In one exercise inspired by the decompositions in Adam et al. (2024), I find that forecasters systematically *underrespond* to the public component of new information, while overreacting to its private counterpart[1]. This pattern rules out diagnostic models, overconfidence and with models of varying attention, which imply overreaction to *any* new information. Relatedly, I include in the testing methodology model's implications that are *unique* to the proposed information structure, validating its suitability and distancing it from the commonly used noisy information framework.

The bias this paper exposes emerges from distinguishing forecasts from conditional expectations, and holds independently of the specified information structure. For instance, anticipating one of the findings, I show that the estimates of information rigidity advanced by Coibion and Gorodnichenko (2015) are effectively lower bounds when considering the counteracting effect of strategic forces. In this case, the test of FIRE is *conservative* with respect to the null of the absence of information frictions. On the other hand, these statistics are used not only for hypotheses testing, but serve also as direct measures of nominal rigidities, such as the degree of price stickiness or the learning ability of economic agents. Consequently, misestimating these metrics has undesirable consequences on our understanding of macroeconomic dynamics, from anticipation effects to the responses to shocks.

Previous studies focused on the rationality of forecasters failed to account for strategic motives in forecasters' objectives, resulting in inference about agents' cognitive abilities that justified the introduction of a wide spectrum of behavioral models. The multitude of viable behavioral models is subject to a classic critique by Sims (1980), famously recast by Sargent (1999) as a "wilderness" of alterna-

[1] For reasons of length and flow, these (and other) findings are not shown in the main body. Appendix O details the empirical specifications and their interpretations.

tives to rational expectations. In this spirit, this study advances a solution by unveiling a new mechanism behind the expectation formation process, explaining the anomalies in a unified rational framework.

Overall, the core insight is that cross-sectional independence across information sets and in the motives of forecasters is an exceedingly restrictive assumption, one that does not appear plausible in the context of professional forecasting. Survey analysts and policymakers interested in gauging "market expectations" need to take that into account when making inference on individual and consensus responses. This paper proposes a unified framework grounded in strategic thinking able to discipline the interpretation of both, and recommends prudence in using surveys as unbiased sources of expectations in macroeconomics and finance.

In summary, the main contributions of this study are four: (*i*) it documents the strategic origins of misspecification in tests of FIRE; (*ii*) it provides evidence that forecasters pursue strategic motivations that go beyond the maximization of accuracy; (*iii*) it proposes a parsimonious model that incorporates public and private information, while contemplating a broad spectrum of strategic incentives in forecasting problem; (*iv*) it formulates a flexible statistical framework that allows to empirically test the model's restrictions, including the ones exclusive to the model proposed.

The remaining sections of the paper are structured as follows: Section 2 describes the data employed and revisits the extant findings on the under and overreaction of consensus and individual forecasts; it also delineates the notion of test misspecification and delves into new empirical observations, highlighting a strategic channel. Section 3 introduces a simple dynamic model of strategic forecasting, elucidating the underlying rationale and drawing a set of predictions. Section 4 aligns these predictions with the empirical observations, targeting a quantification of the strategic motive's direction, and advancing a unified statistical test of the model as a validating exercise. Section 5 concludes.

2 Empirical Observations

2.1 Data

The forecast data I use is sourced from the Survey of Professional Forecasters (SPF). The SPF, a survey managed by the Federal Reserve Bank of Philadelphia, has been operational since 1968 and is the oldest survey of forecasts in the world. I use data up to the pre-pandemic period, i.e., to the last quarter of 2019. Typically, between 40 and 70 anonymous panelists contribute to the SPF. Conducted quarterly, around the end of the second month of the quarter, the SPF offers both average and forecaster-level data. Forecasters are identified by unique IDs and provide predictions for outcomes in the current and subsequent four quarters and more.

The median (mean) duration a panelist contributes is approximately 16 (23) quarters. This results in an unbalanced panel for each variable. Considering the timing of public distribution of the SPF, by the time they make their predictions in quarter t, forecasters are typically aware of the official values of quarterly release variables (e.g., GDP) up to quarter $t-1$ and of monthly release variables (like unemployment rate) up to the previous month.

Table 1: Summary Statistics

Variable	Errors			Revisions		Forecast Dispersion
	Mean	SD	SE	Mean	SD	
Nominal GDP	-0.38	1.95	0.21	-0.15	0.69	1.06
Real GDP	-0.30	1.87	0.20	-0.16	0.60	0.82
GDP Price Deflator	-0.07	1.11	0.14	0.01	0.43	0.62
CPI	-0.26	1.04	0.13	-0.11	0.45	0.53
Real Consumption	0.20	1.59	0.19	-0.06	0.43	0.64
Industrial Production	-1.12	3.90	0.43	-0.32	1.06	1.67
Real Non-Res. Inv.	0.04	5.74	0.78	-0.26	1.74	2.43
Real Res. Inv.	-0.24	8.12	1.13	-0.63	2.33	4.57
Real Fed. Govt Cons.	0.07	3.11	0.39	0.10	1.16	2.24
Real State&Local Govt Cons.	0.04	1.11	0.15	-0.05	0.35	1.02
Housing Start	-3.28	17.98	2.21	-2.04	5.92	8.81
Unemployment	0.06	1.06	0.11	0.04	0.32	0.32
3M Treasury Rate	-0.53	1.15	0.15	-0.19	0.50	0.44
10Y Treasury Rate	-0.52	0.74	0.11	-0.13	0.36	0.39
AAA Corporate bond Rate	-0.50	0.83	0.11	-0.12	0.38	0.51

Notes: All values are in percentages. Columns 1 to 5 show statistics for errors and revisions of consensus (average) forecasts. Errors are calculated as actuals minus forecasts, and actuals are realized outcomes corresponding to the forecasts. Standard errors of forecast errors are calculated using Newey and West's (1994) recommended procedure. Revisions are calculated as forecasts of the outcome made in quarter t minus forecasts of the same outcome made in quarter $t-1$. Columns 6 shows individual-level forecast dispersion, calculated as the mean of quarterly standard deviations of individual-level forecasts. From variable *Nominal GDP* to *Housing Start* the format is the growth rate from the end of quarter $t-1$ to the end of quarter $t+3$. From *Unemployment Rate* to *AAA Corporate Bond Rate* the format is the average level in quarter $t+3$.
Source: SPF, 1968-2019.

Table 1 outlines the variables under study and their main summary statistics at the arbitrary forecast horizon of $t+3$. Macroeconomic variables such as GDP, price indices, consumption, investment, unemployment, government consumption, and financial variables, as T-bills rates and corporate bond yields, are included in the analysis. Most of this study uses as practical application a one-quarter forecast horizon. Data expressed in levels is converted to its implied growth rates[2]. To exclude outliers, and so address the insightful criticisms by Juodis and Kučinskas (2023), I discard forecasts that deviate by more than 5 interquartile ranges from the median value and I only include those forecasters with a minimum of ten observations in the survey[3]. The consensus forecasts are obtained by averaging

[2] Appendix A includes a breakdown of each variable's transformation.

[3] Alternative treatments of outliers deliver almost identical estimates.

individual-level predictions available at any point in time[4].

Revisions in the release of macroeconomic variables are common and influential, as noted in a recent study by Paul (2023). To align the employed data with the forecasters' information set, I focus on initial releases sourced from the Philadelphia Fed's Real-Time Dataset for Macroeconomists. To illustrate, for Nominal Gross Domestic Product (NGDP) growth from quarter $t-1$ to $t+3$, I use the initial release of $NGDP_{t+3}$ in quarter $t+4$ divided by the contemporaneous release of $NGDP_{t-1}$. The practice of aligning the calculated growth rates with the panelists' information sets is common in the literature[5], and does not generally apply to financial variables, which are normally not subject to revisions.

2.2 Consensus underreaction

The empirical analysis this work is based on is the established methodology that leverages survey forecast responses to test the tenets of full information and rational expectations. The pioneering contribution in this realm can be attributed to Coibion and Gorodnichenko (2015) (CG henceforth), which introduced the following class of tests. By drawing on benchmark models of information rigidities – the *sticky information* model by Mankiw and Reis (2002) and the *noisy information* model by Woodford (2001) – CG introduce a regression-based approach to gauge and quantify deviations from the ideal scenario of frictionless information acquisition. More specifically, they focus on the structural relationship between forecast errors and forecast revisions within these models. Below, I present the notation employed in the rest of the paper, by introducing the central regression of the CG methodology[6]:

$$\left(x_{t+h} - \overline{\mathbb{F}}_t x_{t+h}\right) = \alpha + \beta^c \cdot \left(\overline{\mathbb{F}}_t x_{t+h} - \overline{\mathbb{F}}_{t-1} x_{t+h}\right) + u_t \tag{1}$$

Where x_{t+h} represents the actual realization of variable x at time $t+h$ and $\overline{\mathbb{F}}_t x_{t+h}$ is the *consensus* forecast at time t of variable x at horizon $t+h$, defined as the average of individual forecasts: $\overline{\mathbb{F}}_t x_{t+h} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{F}_t^i(x_{t+h})$[7].

Regression (1) captures the degree of inertia in the panelists' forecasting behavior when new information is available. An intuitive interpretation of this exercise is the following: the independent variable (the average revisions) captures the evolution (from $t-1$ to t) of the panelists' information sets, e.g., the arrival of new information at t; when we regress the average forecast errors (dependent variable) on average forecast revisions, we obtain[8] a measure of "how well" on average the consensus is able to process new information when trying to minimize

[4] Appendix A.2 conducts a series of robustness exercises by excluding the early years (1968-1971) of the SPF, which are deemed problematic by some authors.
[5] Appendix A.1 briefly discusses the logic behind the real-time methodology.
[6] The use of $\mathbb{F}$ operator notation is intended to differentiate forecasts from expectations.
[7] Mathematically, $\mathbb{F}_t^i(x_{t+h}) = \arg\max\{\text{Forecaster } i\text{'s optimization problem}\}$
[8] Assuming stable preferences, i.e., the forecasting problem has remained unchanged.

the (mean squared) error of their predictions. If information was perfect and forecasters fully rational (FIRE), it would not be possible to predict future errors using today's revisions, which would be by definition in forecasters' information sets. Hence, under the assumption of FIRE, the implied value for coefficient β^c is $\beta^c = 0$. A positive coefficient β^c would instead suggest that upward revisions today are associated on average with systematic undershooting, a result that has been known in the literature as "consensus underreaction". Reproducing this exercise on a longer sample, I confirm that $\beta^c > 0$ is a robust finding for inflation expectations, a result that generalizes to many macroeconomic variables. Figure 1 displays the coefficients β^c for 15 variables from the Survey of Professional Forecasters.

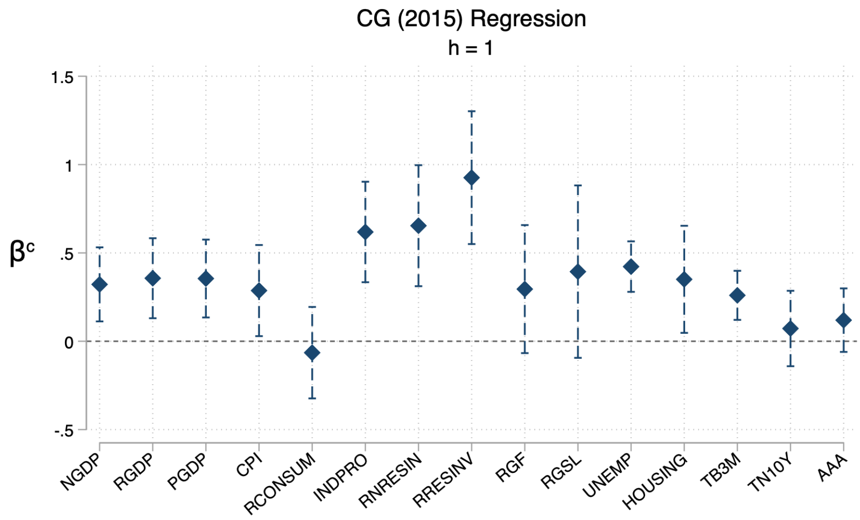

Figure 1: Coefficients from the Coibion and Gorodnichenko (2015) regressions, at a forecast horizon of one quarter; 95% confidence intervals. Variables: *NGDP* (nominal GDP), *RGDP* (real GDP), *PGDP* (price deflator), *CPI* (consumer price index), *RCONSUM* (real consumption), *INDPRO* (industrial production), *RNRESIN* (real nonresidential investment), *RRESINV* (real residential investment), *RGF* (real federal government consumption), *RGSL* (real state and local government consumption), *UNEMP* (unemployment), *HOUSING* (housing starts), *TB3M* (Three-month Treasury rate), *TN10Y* (Ten-year Treasury rate), *AAA* (AAA corporate bond rate).

2.3 Individual overreaction

It is key to emphasize that the restrictions on β^c imposed by the models of information rigidities analyzed in CG – including the null hypothesis of their test – obtain exclusively at the aggregate level, i.e., after averaging across forecasters. In other words, the "map" from the reduced-form estimates to the theoretical objects of interest – the degree of information rigidity – only holds in consensus terms.

Consider the sticky information paradigm as formulated by Mankiw and Reis (2002). In this class of models, at every period forecasters are either drawn to observe the new state ("lucky"), or they do not, implying no revisions to their previous forecasts ("unlucky"). By *averaging* across the two types of forecasters at any given period, and by measuring the degree of predictability that average forecast revisions retain with respect to average forecast errors, the CG test is able to gauge the proportion of lucky/unlucky agents, that can be interpreted as the *stickiness* of information. Similarly, in noisy-information models à la Woodford (2001), all forecasters are exposed to *imprecise* signals of the true underlying process for a variable of interest and extract optimally (*Bayesianly*) the information revealed by weighting the signal based on its perceived precision, making forecast errors unpredictable as a result. The correlation between average errors and revisions arises from a composite effect that stems from the gradual adjustment of forecasters' beliefs to new information. Hence, both relationships obtain exclusively when averaging across agents, and individual regressions per se are uninformative about the presence of information frictions.

It is therefore natural to ask: what do coefficients derived from an individual regression capture? In order to answer, we need to delve into a class of forecasting models in which the rationality of forecasters is put into question. Bordalo et al. (2020) (BGMS henceforth) precisely ventures into this territory, emphasizing the individual nature of the *rational expectations* part of FIRE: at the individual level, the researcher is able to isolate each forecaster's information set and its use in the forecasting problem. By distorting agents' learning process away from the Bayesian paradigm, BGMS advance a model of *diagnostic* expectations, inspired by the seminal work of Kahneman and Tversky (1972). In this model, agents overweigh information that is perceived as more salient. In such a way, BGMS rationalize the (mostly) negative coefficients obtained by estimating the following pooled regression:

$$x_{t+h} - \mathbb{F}_t^i\left(x_{t+h}\right) = \alpha^{p,i} + \beta^p \Big(\mathbb{F}_t^i\left(x_{t+h}\right) - \mathbb{F}_{t-1}^i\left(x_{t+h}\right)\Big) + u_t^{p,i} \tag{2}$$

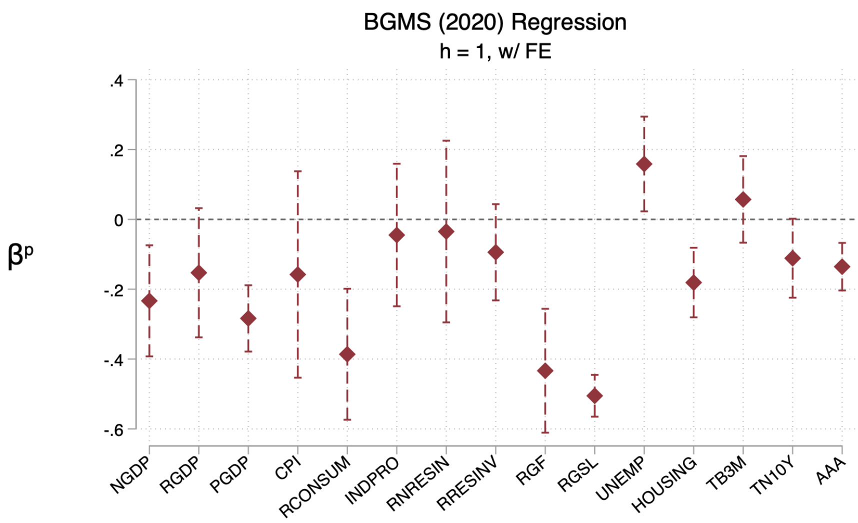

Figure 2: Coefficients (pooled) from the Bordalo et al. (2020) individual regression with forecaster fixed effects. Forecast horizon of one quarter; 95% confidence intervals.

If forecasters were fully rational *and* exclusively interested in the accuracy of their predictions, it should not be possible to predict individual future errors using present information – justifying the null hypothesis of $\beta^p = 0$[9]. As displayed in Figure 2, β^p is negative for most variables, rejecting the null of rationality in the direction of overreaction. In fact, a negative coefficient implies that upwards (downwards) revisions are associated with downwards (upwards) errors, suggesting that when they do, individual forecasters revise *too much*[10].

2.4 Strategic Forecasting

Do agents respond to social incentives when announcing their forecasts? That is, do they take into account the consensus response and their relative positioning with respect to it? Such behavior would imply the introduction of a wedge between their true conditional expectations and their announced survey forecasts.

There exists abundant evidence that forecasters contemplate the consensus when formulating their predictions. However, the direction of the bias is debated and appears to vary with factors such as the institutional context or the area of expertise of forecasters[11]. Studies by Gallo et al. (2002), Lamont (2002), and Hong and Kacperczyk (2010) examine the role of reputational incentives in professional forecasting, highlighting how career concerns lead forecasters to align closely with consensus predictions. Forecasters motivated by job security or advancement are more likely to "herd", especially when incentives diverge between employer and analyst. Conversely, Ottaviani and Sørensen (2006a), Ottaviani and Sørensen (2006b), and Bernhardt et al. (2006) find that, under reputational and reward structures, forecasters favor bold, distinctive predictions ("anti-herding"). Similarly, Laster et al. (1999) shows that forecasters balance accuracy with visibility, reflecting complex strategic incentives. This paper is ex ante agnostic on the direction of the strategic incentive and decouples *forecasts* from true *beliefs* through the lens of a structural model that entertains both possibilities[12].

I add to the empirical side of this literature by proposing as a summary of the data an augmented regression, which augments the individual specification of Bordalo et al. (2020) with a measure of the deviation from consensus (DFC). In the next Section, I will lay out a model in which I interpret these coefficients in a strategic setting where individual forecasters *do not know* the consensus forecasts at the time they announce theirs[13], and Appendix D.1 presents a set of alternative timing specifications using (i) the *past* DFC $\mathbb{F}^i_{t-1}(x_{t+h}) - \overline{\mathbb{F}}_{t-1}(x_{t+h})$; or (ii) a "mixed"

[9] In Appendix C, I provide a formal proof of this statement.

[10] However, forecasters underrespond to the *public* component of new information, which contradicts diagnostic models and reinforces the justification of the strategic angle I advance next. See Appendix O for a discussion.

[11] Appendix B provides a schematic review of the literature, grouped by finding.

[12] Appendix B.5 explains how strategic motives affect also anonymous surveys.

[13] This is a conservative assumption, as evidence suggests participants infer the consensus due to multiple non-synchronous surveys. Participants interviews corroborates the point.

DFC $\mathbb{F}_t^i(x_{t+h}) - \overline{\mathbb{F}}_{t-1}(x_{t+h})$[14]. Findings are robust.

Formally, I run:

$$\underbrace{x_{t+h} - \mathbb{F}_t^i(x_{t+h})}_{\text{Forecast Error}} = \alpha_A^{p,i} + \beta_A^p \Big(\underbrace{\mathbb{F}_t^i(x_{t+h}) - \mathbb{F}_{t-1}^i(x_{t+h})}_{\text{Forecast Revision}} \Big) + \gamma_A^p \Big(\underbrace{\mathbb{F}_t^i(x_{t+h}) - \overline{\mathbb{F}}_t(x_{t+h})}_{\text{Deviation from Consensus}} \Big) + u_{A,t}^{p,i} \quad (3)$$

A few remarks are in order. First, I find that across all variables and horizons, the point estimates of β_A^p in the augmented regression are consistently higher than in the baseline specification. This suggests the presence of a downward bias: if the error term in equation (2) is truly orthogonal to the forecast revision, introducing any other variable to the regression should not systematically alter the estimate of β_A^p. Second, I find that γ_A^p is negative and significantly different from zero for every variable in the sample, which I demonstrate below can be explained by the tendency of analysts to exaggerate their predictions[15]. In fact, forecasters may manipulate their forecasts to stand out, e.g., reporting higher (lower) values when theirs are above (below) the consensus estimate, leading on average to more positive (negative) forecast errors, both scenarios showing up as a negative value for γ_A^p[16]. Figure 3 and 4 display these findings.

A possible (previously mentioned) concern is about the possibility that $\overline{\mathbb{F}}_t(x_{t+h})$ might not be in the information set of agents at time t (a debated assumption, see footnote 13 and a discussion in Appendix D.2, D.3). I address this in two ways: first, by running an alternative specification using the lagged consensus estimate, as proposed by Bordalo et al. (2020)[17]. Second, regardless of the cogency of the robustness checks, I propose a model *conservatively* assuming agents do *not* observe others' contemporaneous forecasts or the consensus. The advantage of the contemporaneous consensus approach is mostly theoretical: in fact, even if (3) does not allow to explicitly test the hypothesis of truthful revelation (see footnote 16), all other specifications do, and after comparing these regressions and observing strong robustness in the results, I chose to keep (3) to represent the closest equivalent to the ideal *unfeasible* regression, which features agents' *contemporaneous expectation* of the consensus estimate (equation (1) in Appendix D.1).

Finally, under the extreme case of FIRE – irrespectively of the chosen forecasters' objective – the signal perfectly reveals the state and such knowledge is translated

[14] Variations to (3) give analogous results. For instance, I discuss the *observability* of the DFC term, or check for asymmetric coefficients of positive/negative DFC in Appendix D.2. Also, anticipating the model, agents' quadratic loss functions will generate a F.O.C. in which the square term decays, making the resulting linear term immediately comparable to (3).

[15] An omitted variable bias (OVB) could be posited. However, this bias arises as a "left-hand side" issue from specifying forecast errors as metric of performance, and not as a "right-hand side" issue, as in common instances of OVB. Therefore, I use the language of "model misspecification".

[16] Appendix D.3 explores why $\gamma_A^p \neq 0$ even when forecasters reveal their expectations truthfully. A lagged-consensus regression ($\gamma_A^p = 0$ under truthful revelation) shows $\hat{\gamma}_A^p < 0$ is robust.

[17] Bordalo et al. (2020), p. 2761: "*In Appendix C, Table C7 we address this mechanism by controlling in the pooled specification of equation (2) for the deviation of the forecast in quarter* $t-1$ *from the consensus* $(x_{t+h|t-1}^i - x_{t+h|t})$."

by all agents in the same exact forecast. Therefore, the distance from consensus (DFC) term would be zero for all i, making it irrelevant to predict forecast errors.

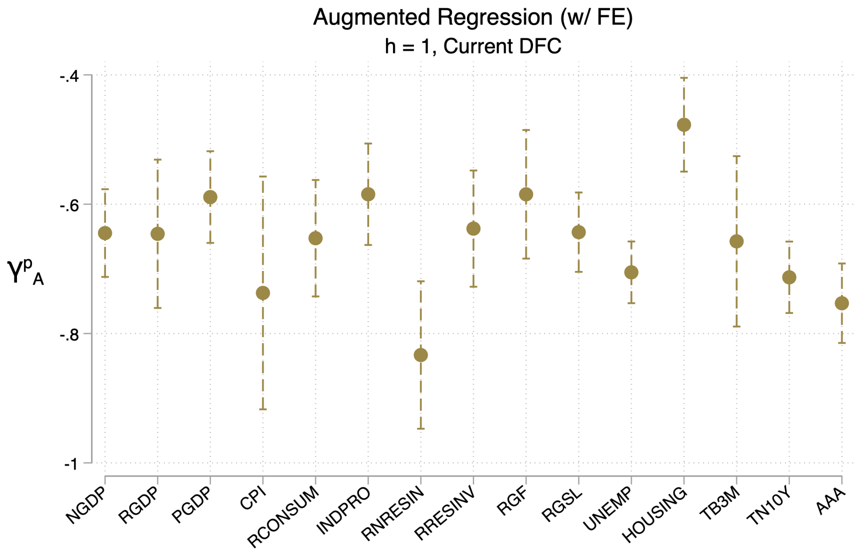

Figure 3: γ^p_A coefficients from the augmented regression (3) with forecaster fixed effects. Forecast horizon of one quarter; 95% confidence intervals.

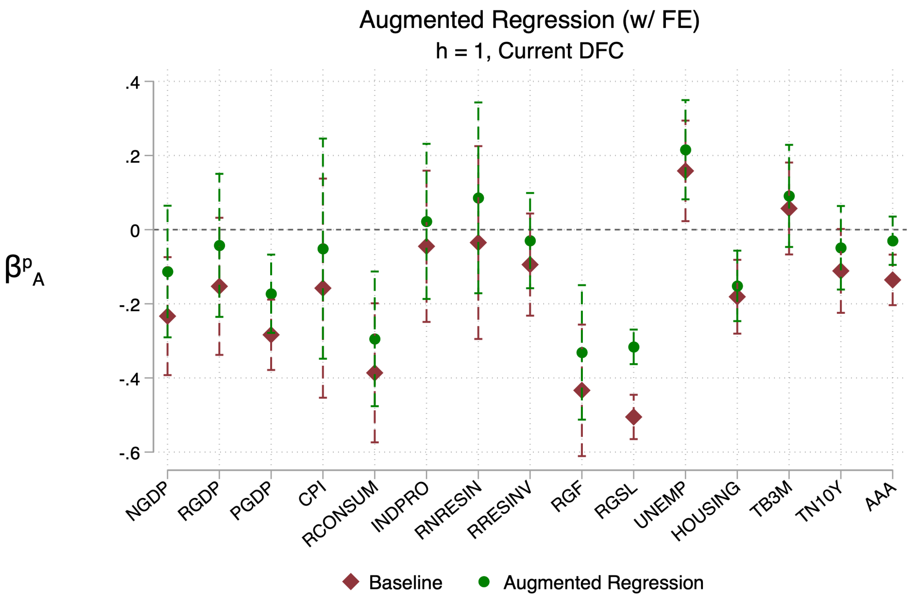

Figure 4: β^p_A comparison between baseline ((2)) and augmented regression (3), both with forecaster fixed effects. Forecast horizon of one quarter; 95% confidence intervals.

In summary, due to the presence of strategic considerations, the estimates of β cannot be mapped either into the degree of information rigidity like CG do, nor in "diagnostic" measures of forecasters' beliefs like BGMS do. Moreover, β ceases to constitute a sufficient statistic to test the null of full information or rationality. Endowed with this result, I advance a model that features (generally distributed) public and private information together with the strategic behavior of forecasters. With its insights, I provide an answer to the question: what can we learn from survey forecast data – which are distorted by strategic incentives – about the use of information by professional forecasters?

3 Model

The model advanced in this study presents elements of innovation with respect to the standard models of information rigidities. The reason for these departures will become evident when the solution is derived. The modeling decisions at the origin of this framework have been inspired by principles of parsimony and simplicity. The idea at the core of the model is to achieve the clearest separation between the *information* and the *strategic* spheres, with the emphasis posed on the latter. Shutting down the information mechanisms in the determination of forecasts, the proposed structure, affine to the affiliated signals used in the auction literature (Milgrom and Weber (1982)), eliminates the entanglement between these two interconnected forces that has made the debate on strategic incentives opaque, and possibly inconclusive. My objective has been to devise the simplest possible model capable of encapsulating the strategic channel in forecasting problems, favoring the interpretability and sharpness of the mechanisms.

3.1 Economic & Information Environment

Forecasters are interested in forecasting variable x, which follows an AR(1) process:

$$x_{t+1} = \rho x_t + u_{t+1} \tag{4}$$

Agents observe the variable of interest (x_t) contemporaneously. That is, the forecasted variable is not *latent* as in most benchmark theories – favoring the introduction of signals *about* x_t –, but is directly observed and incorporated into each forecaster's information set. This modeling choice is motivated by multiple reasons. First, a consideration of intuitive character: many of the results seen revolve around notions of under/overreaction, intended in terms of forecasters' own under/overshooting with respect to the *realization* of the forecasted variable; clearly, these notions postulate the *observability* by the forecast-makers of their own forecasting performance, that is, the contemporaneous observability of x_{t+h} at $t+h$. Secondly, the reduced-form evidence in Section 1 uses the *observed* variable on the left-hand side of each regression. Observing their own forecasting errors is in fact a logic prerequisite for forecasters to even be able to evaluate their performance and learn[18].

I postulate that agents have information about the future realization based on their type τ:

$$u_{t+1} = p_t + \sum_{\tau=1}^{N} y_t^\tau + \epsilon_{t+1} \tag{5}$$

$$p_t \sim i.i.d.\left(0, \sigma_p^2\right) \quad y_t^\tau \sim i.i.d.\left(0, \sigma_y^2\right) \quad \epsilon_{t+1} \sim i.i.d.\left(0, \sigma_\epsilon^2\right) \tag{6}$$

Where y_t^τ is a type-τ specific (private) signal, p_t is a public signal observed by all and ϵ_{t+1} is a pure innovation, not observed by anyone. As I later expand upon, the type-specific signals represent differential exposure to information, while the public signal can be thought of as the portion of information set all participants agree on. Therefore, (5) represents the innovation as the sum of independent (or residualized) components. Note that, differently from almost every other study dealing with similar problems, I do *not* impose Gaussianity nor any other distributional assumption beyond that of independence[19].

The overwhelming majority of papers in this literature impose a Gaussian setting to preserve convenient features of agents' updating behavior. For example, rational inattention frameworks typically use it to yield analytic solutions, and the coordination games à la Morris and Shin (2002) (very relevant for this literature) use Gaussian noise to obtain a unique linear equilibrium. However, Rigos (2022) shows that relaxing the Gaussianity assumption – even in minor ways – produces *qualitatively* different predictions. This model addresses fully such concerns by abstaining entirely from specifying distributional assumptions.

Another element of innovation is in the timing of signals. Endowed with the observation of x_t, the information revealed to forecasters by the signals concerns the *future*. Due to the autoregressive structure of x_{t+1}, this translates the signals into information about tomorrow's innovations, denoted by u_{t+1}[20]. This timing convention fully describes the data generating process underlying to the economic environment, without relying on exogenously imposed structure as in conventional noisy information models. As it will be later expanded upon, this feature allows researchers to disentangle the observational equivalence arising from the isomorphic

[18]There are deep technical reasons for adopting latent models, which merit a paper of their own. In short, latent structures call for the use of filtering algorithms (e.g., the Kalman filter), that employ a recursive relationship between past (*prior*) and present (*posterior*) expectations, establishing a map between empirical observations (e.g., the β of CG regressions) and theoretical counterparts (e.g., the Kalman gain, the Calvo shares). In Appendix E, I elaborate around the latency of information structures.

[19]A discussion of this relevant relaxation is in Rigos (2022), and I complement it through a thorough statistical analysis on the empirical implausibility of Gaussianity in Appendix M. In short, the overwhelming majority of models relies on this assumption. Removing it has serious consequences in terms of identification, interpretability and estimation (e.g., the ubiquitous method of moments becomes nonviable).

[20]The Online Appendix generalizes to multiple periods signals and correlated signals. The insights are unchanged.

effects of the precision of signals and the strength of the strategic channel.

Moreover, it is typical in the literature to model private information as including a random component that has no basis in the underlying economic fundamentals. I advance an alternative formulation based on fundamentals that are observed or not by different agent types. For completeness, I show in Appendix J that this framework can accommodate the restrictions imposed by the standard noisy information model à la Woodford (2001), so that it can be considered a generalization of the benchmark structures of information rigidity[21].

To summarize, the information set of every individual of type τ is:

$$\Omega_t^\tau = \{x_t, p_t, y_t^\tau\} \cup \Omega_{t-1}^\tau \tag{7}$$

There exists N types, with an equal measure of agents for each. *Types* can be interpreted as sources of differential exposure to information. For instance, forecasters in different industries have specialized information[22]. This modeling approach can also capture other dimensions of heterogeneity among forecasters, such as differences in information costs or benefits, as in conventional stories of rational inattention à la Sims (2003) or in the sparsity-based theories of bounded rationality by Gabaix (2014). In rational inattentive models, agents choose information to maximize expected utility, forming posterior beliefs in Bayesian fashion; in sparsity-based frameworks, agents build a simplified (sparse) representation of the world, varied across agents due to differences in preferences, costs, or cognitive abilities.

3.2 Forecaster's Problem

Forecasters of any type face dual incentives when making predictions. First, they aim to ensure their forecasts are accurate. Second, they are influenced by the consensus, either wanting to align with it or differentiate from it. In essence, forecasters are balancing the objective of accuracy with the inclination to either align with or deviate from the consensus.

Agents' preferences are described by the following loss function[23]:

$$\min_{\mathbb{F}_t^\tau(x_{t+h})} \mathcal{L}: \quad \mathbb{E}\left\{\left[\left(x_{t+h} - \mathbb{F}\left(x_{t+h}|\Omega_t^\tau\right)\right)^2 + \phi\left(\bar{\mathbb{F}}_t\left(x_{t+h}\right) - \mathbb{F}\left(x_{t+h}|\Omega_t^\tau\right)\right)^2\right]\middle|\Omega_t^\tau\right\} \tag{8}$$

Where ϕ represents the magnitude of the strategic motive[24], and can be thought

[21] I also precisely characterize the restrictions of the map between the two structures.

[22] The SPF provides a direct example, documenting since 1990 the industry of each survey respondent, grouped (broadly speaking) into finance, non finance and others.

[23] $\mathbb{F}\left(x_{t+h}|\Omega_t^\tau\right)$ represents the announced forecast of x_{t+h} of an agent with information set Ω_t^τ. It is an equivalent but more precise notation for $\mathbb{F}_t^\tau\left(x_{t+h}\right)$, used earlier.

[24] Note that ϕ is common across forecasters to parallel the evidence in Figure 2, where the regression is *pooled* and a unique β^p estimated. Heterogeneity in ϕ can be implemented by allowing (8) to be individual-specific (ϕ^i)–using forecaster-by-forecaster coefficients, as in Appendix F.

as a reduced-form device to generate the relevance of the strategic considerations. In fact, although it is straightforward to microfound these preferences (Appendix B provides a wide array of possibilities), I argue that a reduced-form approach is the more suitable option in this case for two reasons: first, it serves the intended contribution of leaving ϕ without a predetermined sign imposed by the specific mechanisms imposed; second, given the wide heterogeneity of possible microfoundations, it encapsulates the generality of the class of incentives captured by the model. Some remarks follow:

(i) $\mathcal{L}$ features a trade-off between accuracy and distance from consensus[25];

(ii) $\mathcal{L}$ produces *beauty contest* dynamics, that is, generates higher order beliefs; forecasters have in fact the incentive to second guess others' forecasts in order for them to make the optimal decision;

(iii) $\mathcal{L}$ can accommodate strategic complementarity ($\phi > 0$) or substitutability ($\phi < 0$);

(iv) when $\phi = 0$, $\mathcal{L}$ nests the conventional problem of accuracy maximization, i.e., the minimization of mean squared errors.

The first-order condition characterizes the optimal forecast, denoted by $\mathbb{F}^*(x_{t+h})$:

$$\underbrace{\mathbb{F}\left(x_{t+h} \mid \Omega_t^\tau\right)}_{\mathbb{F}^*(x_{t+h})} = \frac{1}{1+\phi}\mathbb{E}\left(x_{t+h}|\Omega_t^\tau\right) + \frac{\phi}{1+\phi}\mathbb{E}\left(\overline{\mathbb{F}}\left(x_{t+h} \mid \Omega_t^\tau\right)\right) \tag{9}$$

$$\hat{x}_{t+h|t}^\tau = \frac{1}{1+\phi}\mu_{t+h|t}^\tau + \frac{\phi}{1+\phi}\mathbb{E}_\tau\left(\bar{\hat{x}}_{t+h|t}\right)$$

where the second line rewrites compactly equation (9)[26].

As expected, the F.O.C. reflects the key forces at work: the optimal forecast of any type-τ agent is determined by a linear combination of her conditional expectation (the true belief, $\mu_{t+h|t}^\tau$) and her expectation of the average contemporaneous forecast, with weights proportional to the strength of the strategic motive, ϕ. Intuitively, a way to read (9) is to interpret the optimal announced forecast as the forecaster's true belief *plus* a wedge introduced by the strategic *bias*. It is immediately apparent that by setting $\phi = 0$ (e.g., "shutting down" the strategic motive), we obtain the familiar result that the optimal forecast for an accuracy-maximizer agent is indeed her conditional expectation.

Furthermore, the first-order condition (9) encapsulates a fixed point relation, that can be solved using a linear conjecture. In order to proceed, we need to solve for the fixed point using a linear guess[27]. Without loss of generality, I will focus on a

[25] Anonymity uncouples survey responses and respondents. I discuss and extend previous results providing evidence it does not constitute a threat to strategic behavior in Appendix B.5.

[26] A discussion of the relative Nash Equilibrium is outlined in Appendix B.4.1.

[27] Another possibility to solve the fixed point problem is the "repeated substitution" technique proposed by Woodford (2001).

one-step-ahead forecast, $h = 1$. The conjecture is formulated as follows

$$\hat{x}^{\tau}_{t+1|t} = \theta_0 x_t + \theta_1 p_t + \theta_2 y^{\tau}_t \tag{10}$$

where θ_0, θ_1 and θ_2 are scalar weights to be determined as functions of exogenous parameters. Then, knowing that

$$\begin{aligned} \mu^{\tau}_{t+1|t} &= \mathbb{E}\left(x_{t+1}\middle|\Omega^{\tau}_t\right) \\ &= \mathbb{E}\left(\rho x_t + p_t + \sum_{\tau=1}^{N} y^{\tau}_t + \epsilon_{t+1}\middle|\Omega^{\tau}_t\right) \\ &= \rho x_t + p_t + y^{\tau}_t \end{aligned} \tag{11}$$

we conclude

$$\begin{aligned} \hat{x}^{\tau}_{t+1|t} &= (1+\phi)^{-1}\left[\rho x_t + p_t + y^{\tau}_t\right] + \phi(1+\phi)^{-1}\mathbb{E}\left[\sum_{\tau=1}^{N}\frac{\hat{x}^{\tau}_{t+1|t}}{N}\middle|\Omega^{\tau}_t\right] \\ &= (1+\phi)^{-1}\left[p x_t + p_t + y^{\tau}_t\right] + \phi(1+\phi)^{-1}\mathbb{E}\left[\theta_0 x_t + \theta_1 p_t + \theta_2\frac{\sum_{\tau=1}^{N} y^{\tau}_t}{N}\middle|\Omega^{\tau}_t\right] \\ &= (1+\phi)^{-1}\left[\rho x_t + p_t + y^{\tau}_t\right] + \phi(1+\phi)^{-1}\left[\theta_0 x_t + \theta_1 p_t + \frac{\theta_2}{N} y^{\tau}_t\right]. \end{aligned}$$

Solving for the coefficients delivers

$$\theta_0 = \rho \quad \theta_1 = 1 \quad \theta_2 = \frac{1}{1+\phi\left(1-N^{-1}\right)}$$

The optimal announced forecast is thus

$$\hat{x}^{\tau}_{t+1|t} = \rho x_t + p_t + \frac{1}{1+\phi\left(1-N^{-1}\right)} y^{\tau}_t \tag{12}$$

which verifies the conjecture. Equation (12) constitutes the solution to the forecaster's problem, that is, her announced forecast. Notice that the solution is consistent with the theoretical understanding of the literature on the weighting of public and private information, as schematized below:

No Strategic Behavior	$\phi = 0$	$\iff$	$\theta_2 = 1$
Strategic Complementarity	$\phi > 0$	$\iff$	$\theta_2 < 1$
Strategic Substitutability	$\phi < 0$	$\iff$	$\theta_2 > 1$

That is, in the absence of strategic incentives, the weights on public and private information are purely determined by their objectively perceived importance in (5), and this does not rely on Gaussian disturbances. On the other hand, in the context of strategic complementarity — where forecasters aim to align their predictions closer to the consensus — agents will rationally assign less weight to their private information. This action aligns with their dual goal: ensuring prediction accuracy and avoiding divergence from the consensus. Symmetrically, a negative ϕ creates a trade-off between precision and publicity, and forecasters pursue such balance by overweighing privately observed signals.

The information structure, based on non-latent processes and a transparent data generating process, allows the *disentanglement* of strategic and information channels, favoring interpretability. Once N is estimated, this structure makes it possible to interpret a higher (lower) reliance on private signals, relative to public signals, as an unambiguous result of an increase (decrease) in ϕ. This interpretation is not confounded by the observationally equivalent effects of signal precision and strategic motives in standard noisy information models. Note that N constitutes the only conceptually new element with respect to these models but should *not* be interpreted as providing an extra degree of freedom "transferring" the information/strategic ambiguity of previous models. However, all estimation results (based on baseline and permutations of multiple targeted moments) deliver estimates of $\hat{N}$ nearly equal to 2[28], validating the value set in the baseline and eliminating the degree of freedom potentially offsetting the model's improvements in terms of interpretability (one could in fact argue that $\hat{\phi}$ depends the number of types).

Comparing the optimal announced forecast $\hat{x}^{\tau}_{t+1|t}$ with the expectation $\mathbb{E}\left(x_{t+1} \mid \Omega^{\tau}_{t}\right)$ as in (11), it is immediate to notice that they differ only in the weight assigned to private information; when $\phi = 0$, we retrieve the identity between expectation and forecast. Furthermore, it is worth noting that the solution incorporates N^{-1} in the denominator. This stands for the general model's scenario with N distinct types of agents. Intuitively, the extreme case where $N = 1$ aligns with a setting where strategic motives are absent, as all forecasters recognize that all accessible information is inherently public.

How should we think about the economic meaning of ϕ? Gallo et al. (2002), Lamont (2002), Hong and Kacperczyk (2010) study what they call *implicit incentive schemes* in professional forecasting. These studies model reputational and career motives that underlie the decisions of analysts through principal-agent problems in which the incentives of the employer (firm) and the analyst (forecaster) are misaligned. They notice that when concerns about job separation or career advancement exist, forecasters tend to become more risk-averse and *herd*, i.e., reduce their perceived distance from the consensus. Similarly, Clement and Tse (2005) documents a negative correlation between herding behavior and several observable dimensions of forecaster characteristics, such as (*i*) forecaster seniority, (*ii*) brokerage firm size and (*iii*) prior forecasting accuracy, thereby uncovering a channel between the probability of being laid off and the tendency of analysts to herd.

On the other hand, a series of studies as Ottaviani and Sørensen (2006a), Ottaviani and Sørensen (2006b) show that under specific circumstances, forecasters display the opposite propensity, that is, they tend to make bold forecasts that

[28] This holds irrespectively of the choice of algorithm employed and of initial points specified (e.g., $N = 1, 10, 100$). For details, please see Appendix I.3

tend to differentiate themselves from their peers. These papers posit two scenarios under which this is both theoretically and empirically found: the so-called (i) *reputational cheap talk* games, where the "reputation" of a privately informed forecaster is updated by the market conditioning on their announcements; (ii) "winner takes all" contests, where the compensation scheme is skewed towards the top of the distribution. Comparable conclusions are presented in an empirical study by Bernhardt et al. (2006), which shows that the conditional probability of a bias toward the consensus is lower than in the opposite direction, implying a tendency of forecasters for *anti-herding*. Laster et al. (1999) obtain an analogous result showing that forecasters' compensation is a function of both accuracy and publicity[29].

3.3 A unified explanation for the empirical evidence

Given a clearly defined solution for the forecaster problem, the next step involves evaluating the model's predictions against the empirical findings. In practice, we are interested in mapping the optimal forecast predicted by the model, (12), to the empirical estimates covered in Section 2.

3.3.1 Predicted values for the individual regression

For the individual case, the regression of reference is (2)[30]. Note that the prediction in (9) implies that the left-hand side of that regression is given by

$$\begin{aligned} FE^{\tau}_{t+1} &= x_{t+1} - \hat{x}^{\tau}_{t+1|t} \\ &= \sum_{\tau' \neq \tau}^{N} y^{\tau'}_t + \varepsilon_{t+1} + \frac{\phi\left(1 - N^{-1}\right)}{1 + \phi\left(1 - N^{-1}\right)} y^{\tau}_t \end{aligned}$$

While the independent variable takes the following form

$$\begin{aligned} FR^{\tau}_t &= \hat{x}^{\tau}_{t+1|t} - \hat{x}^{\tau}_{t+1|t-1} \\ &= \rho \left[\varepsilon_t + \underbrace{\frac{\phi\left(1 - N^{-1}\right)}{1 + \phi\left(1 - N^{-1}\right)}}_{Q} y^{\tau}_{t-1} + \sum_{\tau' \neq \tau}^{N} y^{\tau'}_{t-1} \right] + p_t + \underbrace{\frac{1}{1 + \phi\left(1 - N^{-1}\right)}}_{D} y^{\tau}_t \end{aligned}$$

Then, it follows that β^p is:

$$\beta^p = \frac{\mathrm{Cov}(FE^{\tau}_{t+1}, FR^{\tau}_t)}{\mathrm{Var}(FR^{\tau}_t)}$$

[29] Appendix B.4 shows how strategic incentives can exist in equilibrium even in the presence of rational expectations equilibrium (REE) solution concepts.

[30] Again, I illustrate using the $h = 1$ case.

$$\beta^p = \frac{\text{Cov}\left(\sum_{\tau\prime\neq\tau}^{N} y_t^{\tau'} + \varepsilon_{t+1} + Q y_t^{\tau}, \rho\left[\varepsilon_t + Q y_{t-1}^{\tau} + \sum_{\tau'\neq\tau}^{N} y_{t-1}^{\tau'}\right] + p_t + D y_t^{\tau}\right)}{\text{Var}\left(\rho\left[\varepsilon_t + Q y_{t-1}^{\tau} + \sum_{\tau'\neq\tau}^{N} y_{t-1}^{\tau'}\right] + p_t + D y_t^{\tau}\right)}$$

$$\beta^p = \frac{QD\sigma_y^2}{\sigma_y^2\left[\rho^2 Q^2 + \rho^2(N-1) + D^2\right] + \rho^2\sigma_\varepsilon^2 + \sigma_p^2} \tag{13}$$

We know that when $\phi < 0$ (> 0) forecasters desire to stand out (herd) with respect to the consensus. According to the model's prediction, the sign of β^p depends on the direction of the strategic motive, indeed the sign of ϕ. In particular:

$$\beta^p < 0 \iff \phi < 0 \tag{14}$$

Hence, in order to rationalize the findings of overreaction documented by BGMS, the model suggests that forecasters' responses need to be strategic *substitutes*. Unsurprisingly, the wish to stand out manifests in the rational over-weighting of private signals, y_t^{τ}, which is emphasized as information exclusive to types τ. This mechanism is consistent with a concurrent story discussed in Gemmi and Valchev (2023)[31].

Conceptually, the strategic channel affects the covariance structure of forecast errors and revisions, giving rise to intertemporal dependencies. Note that in the special case when $\phi = 0$, such intertemporal relation dissolves, generating, as expected, a $\beta^p = 0$. That is, the forecast error and the forecast revision are orthogonal objects, as predicted by any Bayesian framework. Intuitively, in a $min\{MSE\}$ context, individuals use (efficiently) all information available at t to maximize their forecasting accuracy. To see this mathematically, notice:

$$FE_{t+1}^{\tau} = \sum_{\tau'\neq\tau}^{N} y_t^{\tau'} + \varepsilon_{t+1}$$

$$FR_t^{\tau} = \rho\sum_{\tau'\neq\tau}^{N} y_{t-1}^{\tau'} + \rho\varepsilon_t + p_t + y_t^{\tau}$$

confirming that $\text{Cov}(FE_{t+1}^{\tau}, FR_t^{\tau}) = 0$ and thus $\beta^p = 0$.

Finally, let me stress that the model allows for the description of the entire data generating process, a feature that will be useful in the estimation stage. For instance, in the individual regression, using just the forecast errors and revisions described above, we can immediately derive a prediction for the residuals and study the variance of such object (imposing, for simplicity, homogeneous variances across types):

$$\text{Var}\left(u_t^p\right) \overset{\sigma_{y^\tau}=\sigma_{y^{\tau'}}}{=} \left\{(N-1) + \left[\frac{\phi\left(1-N^{-1}\right) - \beta^p}{1+\phi\left(1-N^{-1}\right)}\right]^2 + \left(\beta^p\rho\right)^2\left(Q^2 + N - 1\right)\right\}\sigma_y^2 +$$

[31] However, this model does *not* impose strategic substitutability ex ante, accommodating *all* 15 SPF variables, including positive individual coefficients (TB3M and UNEMP).

$$+\beta^{p2}\sigma_p^2 + \left[1+(\beta^p\rho)^2\right]\sigma_\varepsilon^2 \tag{15}$$

3.3.2 Predicted values for the consensus regression

At consensus level, the reference equation is (1). To establish a direct structural correspondence with the objects in (1), equation (12) is averaged across types of information. Given the uniform distribution of forecasters across information types, this procedure allows to retrieve the consensus variables essential for computing the model's predictions regarding the findings in CG. The average announced forecast is:

$$\bar{\hat{x}}_{t+1|t} = \rho x_t + p_t + \frac{\sum_{\tau=1}^{N} y_t^\tau}{1+\phi\left(1-N^{-1}\right)}\frac{1}{N} \tag{16}$$

It is important to recognize that the finiteness of N preserves the significance of private information through the aggregation process. As N approaches infinity, the influence of ϕ diminishes, highlighting the pivotal role of finite N in maintaining the relevance of individual information contributions. Then:

$$\begin{aligned}
\overline{FE}_{t+1} &= x_{t+1} - \bar{\hat{x}}_{t+1|t} \\
&= \underbrace{\frac{(N-1)(1+\phi)}{N+\phi(N-1)}}_{1-\frac{D}{N}}\sum_{\tau=1}^{N} y_t^\tau + \varepsilon_{t+1} \\
\overline{FR}_t &= \bar{\hat{x}}_{t+1|t} - \bar{\hat{x}}_{t+1|t-1} = \bar{\hat{x}}_{t+1|t} - \rho\bar{\hat{x}}_{t|t-1} \\
&= \rho\left[\left(1-\frac{D}{N}\right)\sum_{\tau=1}^{N} y_{t-1}^\tau + \varepsilon_t\right] + p_t + \frac{D}{N}\sum_{\tau=1}^{N} y_t^\tau
\end{aligned}$$

Where, for $\overline{FR}_t$, I have used the fact that $\bar{\hat{x}}_{t+1|t-1} = \rho\bar{\hat{x}}_{t|t-1}$.
Hence, the structural equivalent for the consensus coefficient β^c is:

$$\beta^c = \frac{\text{Cov}\left(\overline{FE}_{t+1}, \overline{FR}_t\right)}{\text{Var}\left(\overline{FR}_t\right)} = \frac{D(N-D)\sigma_y^2}{\rho^2\left[\left(1-\frac{D}{N}\right)^2 N\sigma_y^2 + \sigma_\varepsilon^2\right] + \sigma_p^2 + \frac{D^2}{N}\sigma_y^2} \tag{17}$$

As in the individual case, the coefficients are set-identified for ϕ:

$$\beta^c > 0 \iff \phi > -1 \tag{18}$$

Interestingly, even when $\phi = 0$ (that is, when *forecasts* coincide with *expectations*), the coefficient on forecast revisions is positive. This should not surprise, as even in the absence of strategic motives, the nature of information is incomplete and it is sufficient to generate the kind of belief inertia that a positive β^c implies. More precisely, same-type forecasters will condition on correlated information when making their predictions. In other words, as long as N is finite, private signals do not wash out completely, and their covariance structure preserves in β^c. The role of ϕ is then one of *amplifier/dampener* of the already positive coefficient. Strategic substitutability plays a counteracting force, de facto curtailing the degree of consensus underreaction represented by $\beta^c > 0$: for instance, an extreme level as $\phi = -1$ brings the coefficient to zero, completely offsetting the correlation effect of private information. Symmetrically, the model predicts that the *closer* forecasters wish to be, the more positive β^c will be.

Analyzing conditions (14) and (18), that account respectively for the model restrictions able to accommodate the individual and consensus coefficients, we notice that there exists a region of intersection:

$$\beta^c > 0 \ \wedge \ \beta^p < 0 \iff \phi \in (-1, 0) \tag{19}$$

When this condition holds, the model is compatible with both patterns documented in the literature. That is, a moderate amount of strategic differentiation is able to reconcile the two most studied coefficients in this space, bridging them in a unified rational framework.

3.3.3 Predicted values for the augmented regression

Finally, I proposed the evaluation of an *augmented* regression that incorporates a notion of distance from consensus in the independent variables, as per equation (3). In order to calculate the model-implied coefficients of equation (3), I need a structural expression for the value of DFC_t^τ:

$$DFC_t^i \equiv \hat{x}_{t+1|t}^\tau - \bar{\hat{x}}_{t+1|t} = \frac{N-1}{N+\phi\,(N-1)} y_t^\tau - \frac{D}{N} \sum_{\tau' \neq \tau}^{N} y_t^{\tau'}$$

OLS formulas for the bivariate linear regression provide the direct mapping from the model to the empirical estimates[32]:

$$\beta_A^p = \frac{\text{Cov}(FR_t^\tau, FE_{t+1}^\tau)\,\text{Var}(DFC_t^\tau) - \text{Cov}(DFC_t^\tau, FR_t^\tau)\,\text{Cov}\big(DFC_t^\tau, FE_{t+1}^\tau\big)}{\text{Var}(FR_t^\tau)\,\text{Var}(DFC_t^\tau) - \text{Cov}(FR_t^\tau, DFC_t^\tau)^2} \tag{20}$$

$$\gamma_A^p = \frac{\text{Cov}(DFC_t^\tau, FE_{t+1}^\tau)\,\text{Var}(FR_t^\tau) - \text{Cov}(DFC_t^\tau, FR_t^\tau)\,\text{Cov}\big(FR_t^\tau, FE_{t+1}^\tau\big)}{\text{Var}(FR_t^\tau)\,\text{Var}(DFC_t^\tau) - \text{Cov}(FR_t^\tau, DFC_t^\tau)^2} \tag{21}$$

With these set of facts in mind, I next ask what specification of the model can unifyingly explain the empirical evidence.

[32] In Appendix H, I provide the full closed-form expressions, which – due to the non-null covariance between the regressors – are too convoluted to be displayed here fully.

4 Empirical Application

4.1 Model Identification

In the previous section, I collected a series of predictions that confront the theoretical implications of the model with the key facts emerged from the empirical results. Each prediction imposes certain restrictions on the parameters of the model. Equation (19) summarizes the intersection of two such restrictions able to offer a rationale for the controversial finding of individual overreaction and consensus underreaction. The purpose of this section is to generalize this effort systematically, comprehensively explaining all the presented evidence.

Then, in order to align the predictions with the empirical findings, it is necessary to identify several parameters. Specifically, the model includes the following: N, representing the number of *types* or distinct sources of information; ρ, the AR(1) coefficient, capturing the persistence of the forecasted process; the variances $\sigma^2_{y\tau}$, σ^2_p, and σ^2_ε which correspond respectively to the variability of the private signals of type τ, of the public signal, and of the purely unpredictable component of future innovation; and lastly ϕ, measuring the strength of the strategic motive in forecasters' preferences. For simplicity, the baseline specification considers homogeneous variances across types τ, but this assumption can be simply relaxed.

The selection of empirical moments to identify the model's parameters range from a simple AR(1) fit, to the three key regressions (1), (2), (3) discussed in Section 2. In particular, here I present the results based on an identification approach tailored to target β^p for (2), β^c for (1), along with the augmented coefficients β^p_A and γ^p_A from (3) and the variance of the residuals of the individual regression, $\sigma^2_{u^p}$. Furthermore, the model is general enough to allow the researcher to test virtually any prediction. In the next Section, I include one unique feature of the estimated news representation of information, that is, the variance of the shock to the fundamental process, $\sigma^2_{u_t}$.

Formally, I rewrite the three regressions in deviations from the time average[33] to mirror the fixed effect estimation implemented empirically, in line with the within transformation:

$$\widetilde{FE}^i_{t+h} = \beta^p \cdot \widetilde{FR}^i_t + \tilde{u}^{p,i}_t$$
$$\widetilde{\overline{FE}}_{t+h} = \beta_{CG} \cdot \widetilde{\overline{FR}}_t + \tilde{u}_t$$
$$\widetilde{FE}^i_{t+h} = \beta^p_A \widetilde{FR}^i_t + \gamma^p_A \widetilde{DFC}^i_t + \tilde{u}^{p,i}_{A,t}$$

Define

$$\pi = \left(\beta^p_{BGMS}, \beta_{CG}, \beta^p_A, \gamma^p_A, \sigma^2_{u^p}\right)'$$

[33] Mathematically: $\forall x^i_t : \widetilde{x}^i_t = x^i_t - \frac{1}{T} x^i_t$

$$q_t(\pi) = \begin{bmatrix} \sum_{i=1}^{I_t} \left(\widetilde{FE}_{t+h}^i - \beta_{BGS}^p \widetilde{FR}_t^i \right) \widetilde{FR}_t^i \\ \left(\overline{\widetilde{FE}}_{t+h} - \beta_{CG} \overline{\widetilde{FR}}_t \right) \overline{\widetilde{FR}}_t \\ \sum_{i=1}^{I_t} \left(\widetilde{FE}_{t+h}^i - \beta_A^p \widetilde{FR}_t^i - \gamma_A^p \widetilde{DFC}_t^i \right) \widetilde{FR}_t^i \\ \sum_{i=1}^{I_t} \left(\widetilde{FE}_{t+h}^i - \beta_A^p \widetilde{FR}_t^i - \gamma_A^p \widetilde{DFC}_t^i \right) \widetilde{DFC}_t^i \\ \sum_{i=1}^{I_t} \left[\left(\widetilde{FE}_{t+h}^i - \beta_{BGMS}^p \widetilde{FR}_t^i \right)^2 - \sigma_{u^p}^2 \right] \end{bmatrix} \tag{22}$$

Where I_t denotes the number of individuals responding to the survey in time t. The estimated $\hat{\pi}$ represent the five moments targeted, and are characterized by $T^{-1} \sum_{t=1}^{T} q_t(\hat{\pi}) = 0$.

The parameters of the structural model are given by $\theta = \left(\phi, \sigma_y^2, \sigma_p^2, \sigma_\varepsilon^2\right)'$. Given a value of θ and some separately estimated N[34], equations (13), (17), (62), (63) and (15) imply a predicted value for π. We can estimate the value of θ using minimum distance estimation as reviewed by Chamberlain (1984). This approach shares several features with GMM estimation (Angrist and Newey (1991), Wooldridge (2010)), including the use of empirical moments to serve as targets in the numerical algorithm. Specifically, I estimate θ by minimizing

$$[\hat{\pi} - \pi(\theta)]' \left(T^{-1}\hat{S}\right)^{-1} [\hat{\pi} - \pi(\theta)]$$
$$\hat{S} = T^{-1} \sum_{t=1}^{T} q_t(\hat{\pi}) q_t(\hat{\pi})'.$$

Finally, for the baseline specification N is set to 2, representing two types of private signals. Such choice can be thought of, for instance, in terms of the industry categorization the SPF provides, where forecasters are classified as belonging to broadly-defined categories such as finance or research. Moreover, note the conceptual significance of two types of information sources, representing the *minimum* degree of heterogeneity sufficient to generate the strategic mechanisms at the core of the model. In a more general application[35], I control the robustness and plausibility of the two-types assumption by re-estimating the model including N in the estimands. $N = 2$ turns out to be indeed a very good approximation, and represents the closest integer of the estimated number of types for all variables. Besides, its estimates demonstrate it does not merely shift the strategic/information ambiguity of noisy information frameworks onto another parameter.

Estimating separately $\hat{\rho}$, following conventional practice, we are left with four parameters to be aligned with five moments, thus rendering the system overidentified. Appendix I provides several alternatives to the baseline specification, employing

[34] As discussed, including N in the joint estimands is a possibility I explore in Appendix I.3. In fact, the estimation of N confirms the integer value (2) set in the baseline, and it is robust to a series of numerical techniques and starting values.

[35] Summarized in Appendix I.3, Table 6; see also p. 19 for more insights on this.

other data moments from the empirical regressions, and varying horizons.

Solving the system as specified above provides a good approximation with respect to the designated data moments. Furthermore, the remarkable variability across variables results in a broad spectrum of identified parameters. This aligns with the extensive differences in information and strategic incentives observed across macroeconomic indicators.

Looking at the results for the critical parameter, ϕ, the model suggests an average value for ϕ of -0.51 and a median value of -0.67 across variables. This implies that strategic considerations might account for more than *half* of the weight placed on accuracy in the utility functions of professional forecasters, certainly a non-negligible amount. When examining various specifications, the estimates show expected patterns. The average ϕ value remains roughly constant across forecast horizons. Macroeconomic variables tend to show slightly heightened strategic motivations in short-term forecasts, whereas financial indicators lean more towards strategic considerations in longer-term forecasts. This pattern aligns with the uncertainty associated with the data releases of these variables. Table 2 reports the identified parameters in the baseline specification described above.

Table 2: Identified model parameters

Variable	ϕ	$\frac{\sigma_y}{\sum\sigma}$	$\frac{\sigma_p}{\sum\sigma}$	$\frac{\sigma_\varepsilon}{\sum\sigma}$
NGDP	-0.69868	0.99588	0.00206	0.00206
RGDP	-0.61649	0.80452	0.09165	0.10383
PGDP	-0.70130	0.99751	0.00125	0.00125
CPI	-0.52822	0.75251	0.12606	0.12142
RCONSUM	-1.00000	0.99220	0.00334	0.00445
INDPROD	-0.34834	0.99791	0.00089	0.00119
RNRESIN	-0.19820	0.99793	0.00092	0.00115
RRESINV	-0.17552	0.99898	0.00044	0.00058
RGF	-0.87267	0.99618	0.00109	0.00273
RGSL	-0.89862	0.99720	0.00093	0.00187
UNEMP	0.59707	0.71241	0.15088	0.13671
HOUSING	-0.73690	0.99500	0.00188	0.00313
tb3m	-0.09469	0.46861	0.34492	0.18647
tn10y	-0.74206	0.35916	0.46819	0.17265
AAA	-0.67127	0.44020	0.36540	0.19440

Notes: identification achieved targeting β^p, β^c, β^p_A, γ^p_A, $\sigma^2_{u_p}$. Depicted, $N = 2$, $h = 1$.

As mentioned, the results shown for parameter ϕ can be considered a lower bound for this class of models, due to the conservative choice of N. As shown in Appendix I, alternative specifications show the positive dependency between the degree of information heterogeneity and the estimates of strategic reasoning. Intuitively, such feature reflects the notion that, in order to stand out, forecasters rationally overweigh more private information the more different their peers' priors are. On the other hand, it is important to emphasize that the model was developed to isolate the strategic forces at play, admittedly neglecting the other possible mechanism that might work along the same direction. Then, it is fair to say that ϕ captures also other dimensions of heterogeneity across variables beyond strategy[36].

In what follows, I present some figures comparing the empirical findings of Section 2 and the theoretical predictions of the identified model. Figure 5 displays the β^p coefficients from the individual regression and the model-implied coefficients given the identified parametrization. Figure 6 show the coefficients for the consensus regression. Figure 7 and Figure 8 supplement with the coefficients from the augmented regression. As in the rest of the paper, the forecast horizon is set at $h = 1$, with Appendix I containing additional robustness checks for different horizons. Across variables, the model is able to accommodate the individual behavior of forecasters. Finally, Table 3 summarizes the performance of the model.

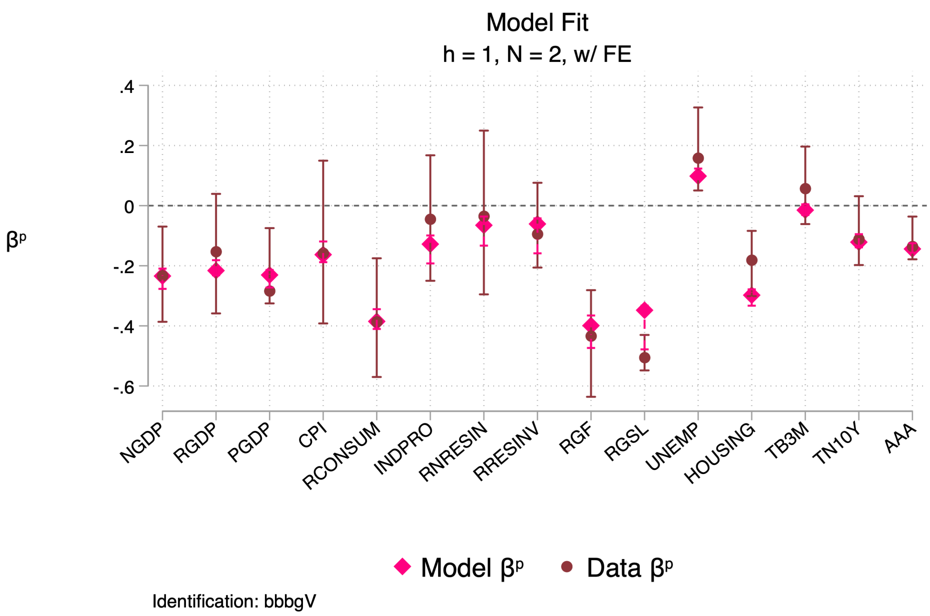

Figure 5: Comparison between β^p from regression (2) at horizon 1, with fixed effects, and the model's predictions under the identified parametrization. Confidence intervals at 95% levels.

[36]Bianchi et al. (2022) and Eva and Winkler (2023) make an insightful argument against the presence of the large biases like the substantial strategic forces I report. Their point focuses on out-of-sample performance, and affects all of the literature dealing with FIRE tests. I address these and similar critiques in Appendix N.

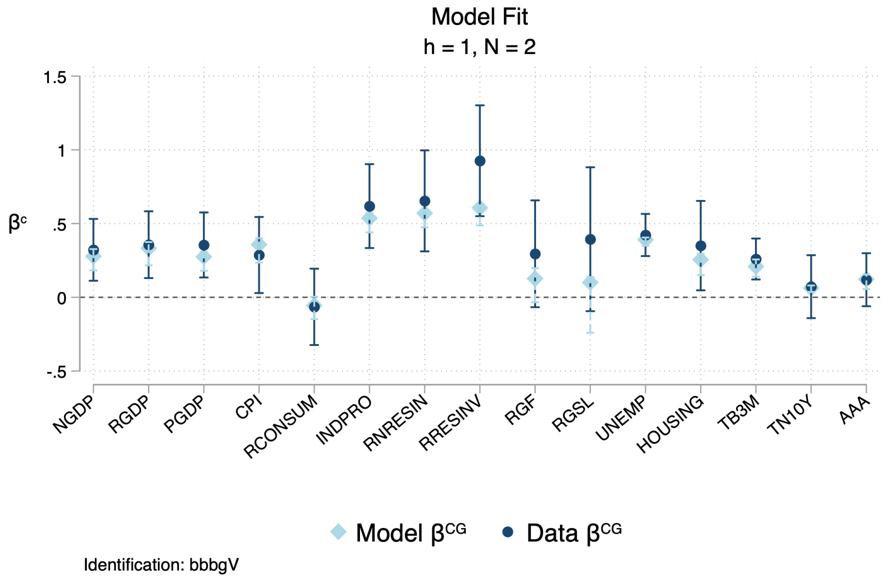

Figure 6: Comparison between β^c from regression (1) at horizon 1 and the model's predictions under the identified parametrization. Confidence intervals at 95% levels.

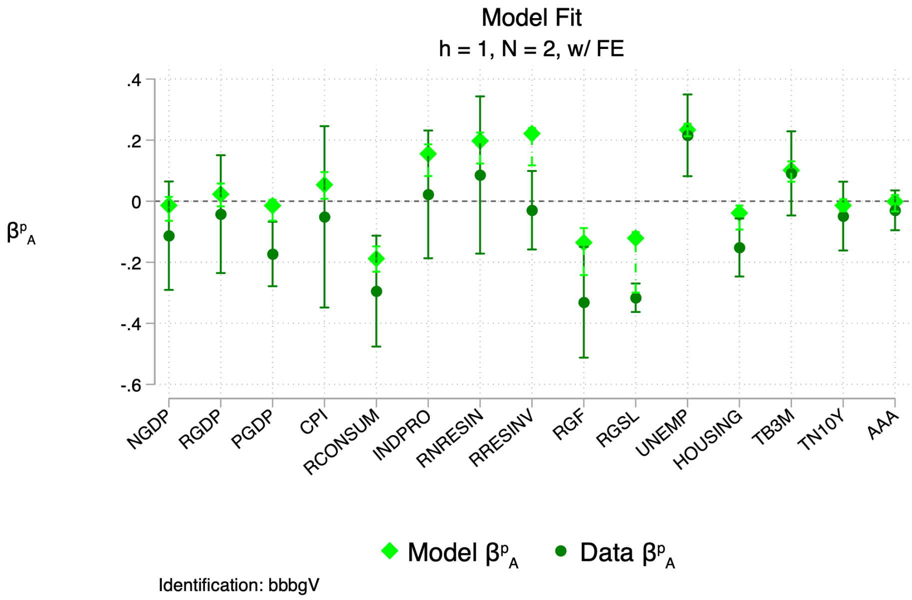

Figure 7: Comparison between β^p_A from regression (3) at horizon 1, with fixed effects, and the model's predictions under the identified parametrization. Confidence intervals at 95% levels.

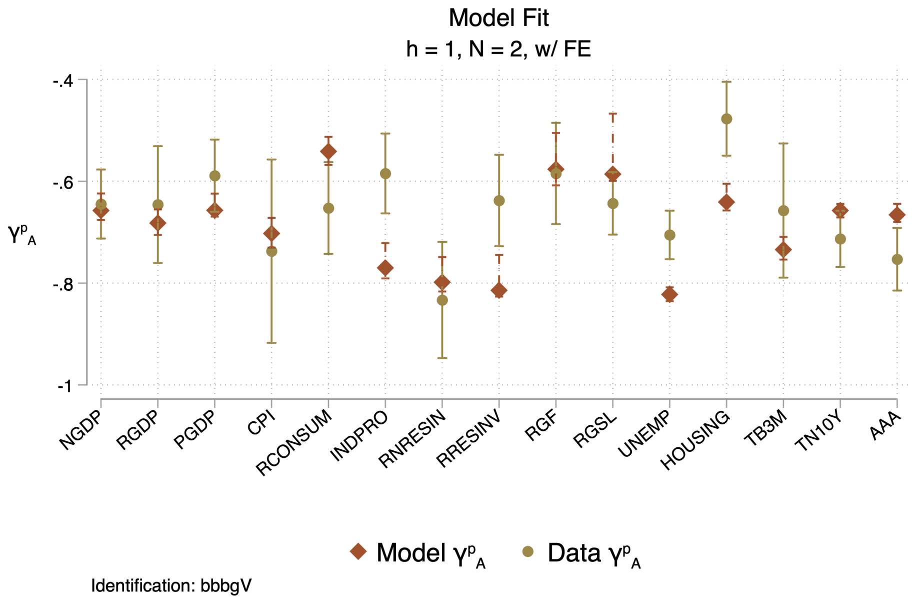

Figure 8: Comparison between γ_A^p from regression (3) at horizon 1, with fixed effects, and the model's predictions under the identified parametrization. Confidence intervals at 95% levels.

Figures 5–8 compare the estimated coefficients with the model's predictions and show an overall good fit, highlighting that the strategic framework replicates the sign and magnitude of most parameters across regressions. In a very few cases (e.g., RRESINV), the model's predictions differ mildly from the data, potentially reflecting measurement issues or omitted factors specific to those series. Nonetheless, the discrepancies remain modest. Table 3 corroborates these visual impressions, indicating strong alignment between empirical estimates and theoretical predictions for the bulk of variables, reinforcing the view that the identified model effectively captures the key dynamics.

Table 3: Model Performance

Variable	β^p		β^c		β^p_A		γ^p_A	
	Model	Data	Model	Data	Model	Data	Model	Data
NGDP	-0.234	-0.229	0.278	0.318	-0.014	-0.102	-0.657	-0.674
RGDP	-0.216	-0.166	0.333	0.349	0.023	-0.042	-0.682	-0.678
PGDP	-0.231	-0.294	0.275	0.346	-0.015	-0.090	-0.657	-0.660
CPI	-0.163	-0.168	0.358	0.357	0.054	-0.017	-0.703	-0.819
RCONSUM	-0.385	-0.390	-0.055	-0.058	-0.188	-0.270	-0.541	-0.679
INDPROD	-0.128	-0.053	0.537	0.681	0.155	0.024	-0.770	-0.601
RNRESIN	-0.065	-0.058	0.570	0.654	0.198	0.079	-0.798	-0.851
RRESINV	-0.061	-0.123	0.607	0.995	0.221	0.010	-0.814	-0.668
RGF	-0.399	-0.445	0.127	0.266	-0.136	-0.346	-0.576	-0.627
RGSL	-0.348	-0.523	0.101	0.234	-0.121	-0.306	-0.586	-0.693
UNEMP	0.098	0.150	0.391	0.421	0.234	0.233	-0.822	-0.695
HOUSING	-0.298	-0.194	0.255	0.331	-0.039	-0.169	-0.641	-0.529
tb3m	-0.015	0.032	0.209	0.244	0.101	0.078	-0.734	-0.637
tn10y	-0.121	-0.119	0.063	0.065	-0.014	-0.046	-0.657	-0.696
AAA	-0.144	-0.149	0.123	0.120	-0.001	-0.034	-0.666	-0.729

Note: identification achieved targeting β^p, β^c, β^p_A, γ^p_A. Depicted, the results for $N = 2$, $h = 1$.

The model I put forward has as cornerstone the strategic behavior of forecasters. By contrast, it is insightful to contemplate the natural null hypothesis that strategy might not indeed be a relevant factor. Figure 9 shows the theoretical predictions when there is no distinction between forecasts and expectations, i.e., the case for $\phi = 0$. Unsurprisingly, in this case forecasters act consistently with the Bayesian paradigm, which I have argued[37] implies that all available information in incorporated in forecast revisions, ruling out predictability. In fact, under such restriction, the model is unable to match the empirical results, and we obtain the expected result of $\beta^p = 0$ for all variables.

Notice, however, that this needs not be the case for the consensus regression. As touched upon in the previous section, the model predicts that the consensus coefficient is positive even in the absence of a strategic factor. This aligns fully with the proposed framework, which features imperfect and incomplete information shared among forecasters of the same type. Moreover, this represents a novel mechanism that abstracts from the *noisiness* of the observed information and the prevalent stories of average underreaction, whose dynamics rely on the idea that forecasters are unable to distinguish true innovations from noise, thus responding prudently in their revisions and generating learning inertia[38]. In this framework instead, agents are fully aware that they possess *true* yet *partial* knowledge about tomorrow's realization x_{t+1}, and engage in a game where each player balances the

[37] and proved, see Appendix C.

[38] A review of similar mechanisms can be found in Angeletos and Lian (2016).

incentive for accuracy with the desire of standing out from the crowd. The role of ϕ, in these scenarios, is one of amplifier or dampener of the already existent effects implied by the learning technology. That is, the finding that $-1 < \phi < 0$, i.e., the strategic overweighting of private information, *contrasts* the intrinsic correlation that generates a positive β^c, without whom we would observe even larger deviations from the null. It is in this sense that the estimates offered by CG can be considered lower bounds of information rigidity.

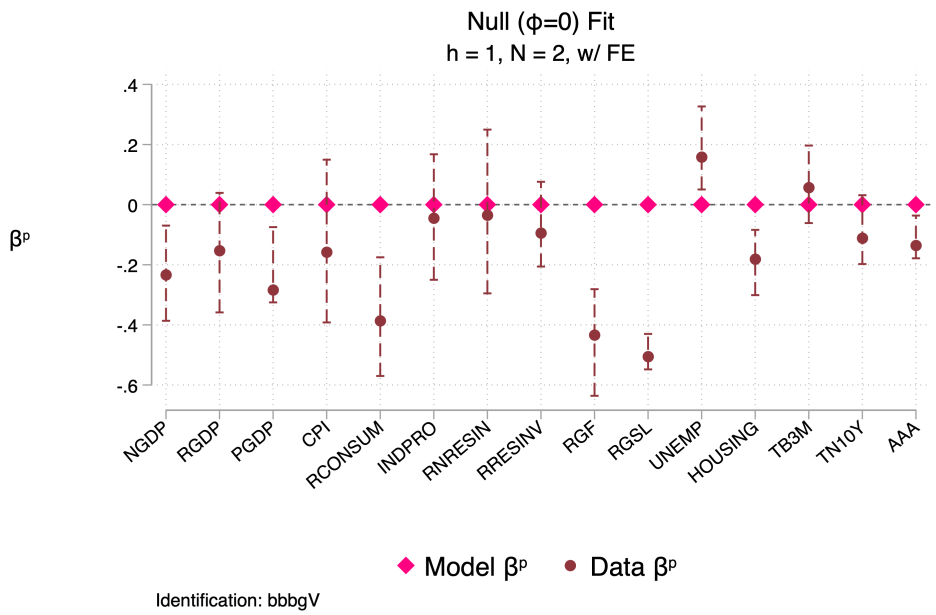

Figure 9: Comparison between β^p from regression (2) at horizon 1, with fixed effects, and the model's predictions under the null of no strategic forces. Confidence intervals at 95% level.

This paper instead advocates for the strategic manipulation of information and does not rely on either noise nor on ad-hoc behavioral features of forecast production. The only indispensable element is the existence of differentiated information, that, as it has been noticed, can be conceptualized in a number of ways, from rational inattentive agents with heterogeneous payoff structures to "sparsely" attentive bounded rational agents.

4.2 Testing the model

While it can be an insightful exercise, comparing the full-fledged strategic model to its natural non-strategic alternative is not a formal test able to validate statistically the goodness of fit of the model. With this in mind, this section advances a unified statistical methodology to evaluate quantitatively the model's performance.
As stressed in the previous section, identification is achieved by targeting a number of moment conditions that exceeds the number of structural parameters to

be estimated. Therefore, the model is overidentified. The estimation algorithm employed in this paper is the Minimum Distance estimator, which belongs to the family of GMM-type estimators that have known statistical properties for hypotheses testing. In what follows, I summarize the implementation of a Chi-square test for overidentifying restrictions inspired by the seminal work of Hansen (1982). Appendix K describes in detail the procedure. Note that, in addition to the parameters from the previous section, I can virtually test *any* restriction that the identified model imposes. To leverage on the unique implications of the information structure I presented, I further add to the tested overidentifying restrictions the model implied variance of u_t:

$$\sigma^2_{u_t} = \sigma^2_{p_t} + N\sigma^2_{y^\tau_t} + \sigma^2_{\varepsilon_y} \tag{23}$$

The conceptual significance of this addition lies in the validation of the adopted news representation, which enables the researcher to describe the forecasting problem organically to the economic environment – i.e., without specifying purely exogenous constructs such as auxiliary noise processes.

Developing equations (22) along the cross-sectional dimension: $\sum_{i=1}^{I_t} \widetilde{x}^i_t = x^{Tot}_t$, I can rewrite the moment conditions as:

$$\hat{\beta}^p \qquad \sum_{t=1}^{T} \left[\left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p \widetilde{FR}^{Tot}_t \right) \widetilde{FR}^{Tot}_t \right] = 0 \tag{24}$$

$$\hat{\beta}^c \qquad \sum_{t=1}^{T} \left[\left(\widetilde{\widetilde{FE}}_{t+h} - \hat{\beta}^c \widetilde{\widetilde{FR}}_t \right) \widetilde{\widetilde{FR}}_t \right] = 0 \tag{25}$$

$$\hat{\beta}^p_A \qquad \sum_{t=1}^{T} \left[\left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p_A \widetilde{FR}^{Tot}_t - \hat{\gamma}^p_A \widetilde{DFC}^{Tot}_t \right) \widetilde{FR}^{Tot}_t \right] = 0 \tag{26}$$

$$\hat{\gamma}^p_A \qquad \sum_{t=1}^{T} \left[\left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p_A \widetilde{FR}^{Tot}_t - \hat{\gamma}^p_A \widetilde{DFC}^{Tot}_t \right) \widetilde{DFC}^{Tot}_t \right] = 0 \tag{27}$$

$$\hat{\sigma}^2_{u^p} \qquad \frac{1}{\sum_{t=1}^{T} I_t - 1} \sum_{t=1}^{T} \left[\left(\hat{\tilde{u}}^{p,2}_t \right)^{Tot} \right] - \hat{\sigma}^2_{\tilde{u}^p} = 0 \tag{28}$$

Define h_t the vector-valued function that stacks the t-evaluated conditions (23)-(28):

$$h_t(\hat{\pi}) = \begin{bmatrix} \left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p \widetilde{FR}^{Tot}_t\right) \widetilde{FR}^{Tot}_t \\ \left(\widetilde{FE}_{t+h} - \hat{\beta}^c \widetilde{FR}_t\right) \widetilde{FR}_t \\ \left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p_A \widetilde{FR}^{Tot}_t - \hat{\gamma}^p_A \widetilde{DFC}^{Tot}_t\right) \widetilde{FR}^{Tot}_t \\ \left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p_A \widetilde{FR}^{Tot}_t - \hat{\gamma}^p_A \widetilde{DFC}^{Tot}_t\right) \widetilde{DFC}^{Tot}_t \\ \widetilde{FE}^{2,Tot}_{t+h} + \hat{\beta}^{p,2} \widetilde{FR}^{2,Tot}_t - 2\beta^p (\widetilde{FE}_{t+h} \widetilde{FR}_t)^{Tot} - (I_t - 1)\hat{\sigma}^2_{\tilde{u}^p} \\ \left(x_{t+h} - \hat{\rho}\, x_{t+h-1}\right)^2 - \hat{\sigma}^2_u \end{bmatrix} \tag{29}$$

Then, one can estimate the asymptotic variance of $h_t(\hat{\pi})$ by[39]:

$$\hat{S} = T^{-1} \sum_{t=1}^{T} h_t(\hat{\pi}) h_t(\hat{\pi})' \tag{30}$$

Its inverse serves as weighting matrix for the test statistic of the Chi-square test for overidentifying restrictions. The test statistics is distributed as follows:

$$\omega = [\hat{\pi} - \pi(\hat{\theta})]' \left(T^{-1}\hat{S}\right)^{-1} [\hat{\pi} - \pi(\hat{\theta})] \sim \chi^2_{n-k} \tag{31}$$

Table 4: Chi-Square Test for Overidentifying Restrictions, $n = 6, k = 4$

Variable	$\omega_{h=1}$	$\omega_{h=2}$	$\omega_{h=3}$	$\chi^2_{0.01,2}$	Rej. $H_0^{h=1}$	Rej. $H_0^{h=2}$	Rej. $H_0^{h=3}$
NGDP	2.317	4.310	2.368	9.210	NO	NO	NO
RGDP	1.491	1.756	1.471	9.210	NO	NO	NO
PGDP	24.586	152.660	307.123	9.210	YES	YES	YES
CPI	14.376	6.702	4.775	9.210	YES	NO	NO
RCONSUM	20.975	12.513	2.438	9.210	YES	YES	NO
INDPROD	1.343	1.181	1.086	9.210	NO	NO	NO
RNRESIN	0.597	1.875	1.140	9.210	NO	NO	NO
RRESINV	0.563	0.540	0.370	9.210	NO	NO	NO
RGF	0.936	0.161	2.300	9.210	NO	NO	NO
RGSL	157.515	37.492	59.926	9.210	YES	YES	YES
UNEMP	0.208	0.106	0.058	9.210	NO	NO	NO
HOUSING	4.304	2.839	3.606	9.210	NO	NO	NO
tb3m	0.069	0.053	0.081	9.210	NO	NO	NO
tn10y	1.198	1.050	1.254	9.210	NO	NO	NO
AAA	0.984	0.893	0.889	9.210	NO	NO	NO

We reject the null if $\omega > \chi^2_{\alpha,n-k}$, where α is the specified level of statistical significance, which is set at 1%. Implementing the test for all fifteen variables in the sample at horizons $h = 1, 2, 3$, I fail to reject the model for 36 out of 45 total

[39] The degree of autocorrelation is low, so the use of Newey-West (1987) is second-order. Appendix K includes details on the autocorrelation coefficients and the autocovariances at lags 1, 2, 3. See Hamilton (1994) for theoretical insights.

specifications[40]. Table 4 summarizes the results of the test.

As in the case of the identification, I include in Appendix K alternative strategies to the baseline case displayed below, including higher orders of over-identification ranging from 5 to 8 targeted moment conditions, and allowing for different degrees of freedom. The results are qualitatively unchanged[41].

4.3 Final Remark: Revisiting the tests of FIRE

This paper's main focus concerned the *measurability* of expectations, by definition unobservable objects. An early conceptualization of this notion comes from Prescott (1977): "Like utility, expectations are not observed, and surveys cannot be used to test the rational expectations hypothesis. One can only test if some theory, whether it incorporates rational expectations or, for the matter, irrational expectations, is or is not consistent with observations". This work operationalizes this idea.

Researchers have implicitly made a crucial assumption that warrants the ability to consider forecasts as expectations. That is, that forecasters aim exclusively at maximizing their performance accuracy. Under this assumption, using data from the Survey of Professional Forecasters (or other forecast surveys, like the Michigan Survey of Consumers, the Livingstone Survey or the Blue Chip Data), provides a valid proxy to measure the predictability of expectations, so that their properties can be informative on the expectation formation processes of responders. This paper challenged this notion and argued that different incentives impact in the determination of the reported empirical regularities.

To summarize, let me point out why this would constitute an issue for the understanding we have on matters of information frictions and rationality. Abstracting from the distinction between individual and consensus forecasts, tests à la CG have a general structure of the following type:

$$\mathrm{FE}_{t+h} = \alpha + \beta \cdot \mathrm{FR}_t + \epsilon_t \qquad \mathrm{H}_0 : \beta = 0 \tag{32}$$

It has been noticed in Section 1 how the null hypothesis $\beta = 0$ rejects distinct notions depending on the level of analysis. To reiterate, while at the consensus level $\beta = 0$ represents a joint hypothesis of full information and rational expectations, at the individual level the focus is restricted on the rationality of forecasters, i.e., their optimal or suboptimal use of information. Optimality is always defined in relation to an objective function, which has ubiquitously specified as a function of forecast errors. Then, the left-hand side of (32) can be interpreted as capturing this feature of the forecasting problem: the objective. On the other hand, the independent variable is the forecast revision, i.e., the *update* induced by the arrival

[40] The model performs less effectively for some of the "inflation" variables, consistently with the difficult forecasting nature of these variables. Moreover, these happen to be cases where forecasters pay special attention to common signals (e.g., oil prices).

[41] Several variations not in the Appendix exist, mostly featuring even stronger validations.

of new information. Hence, revisions are the outcome of the newly arrived signals *and* the specified optimizing behavior of forecasters. Then, it is intuitive to notice how a structure like (32) is unsuited to study problems that abstract from the mere maximization of accuracy.

More formally, this notion translates in an estimation issue. In order to estimate unbiasedly and consistently β, exogeneity is required ($\mathbb{E}\left(\epsilon_t \mid \mathrm{FR}_t\right) = 0$). In this paper, I have shown that this condition fails to hold empirically because of the presence of strategic incentives in forecasters' responses.

5 Conclusions

In this paper, I set out to investigate the implications of using forecasts as proxies for expectations in tests of Full Information Rational Expectations (FIRE). I have addressed this question and provided a consistent rationale for the apparent anomalies of under/overreaction observed in influential contributions to the behavioral macroeconomics literature over the past decade.

Forecasts do not always align with expectations, and the conditions under which they do are often restrictive and empirically implausible in the context of professional forecasting. This discrepancy implies that conventional tests might be overlooking, at a minimum, one critical dimension of optimization, rooted in strategic behavior.

Baseline rational models fail to account for instances where the individual coefficients β^p are not zero. Introducing strategic reasoning into a model of forecasting can, on the contrary, explain this robust feature of the data, without resorting to behavioral theories, such as salience, diagnostic expectations, experience, asymmetric attention, and many others that echo the common saying in the profession "outside of rationality, *anything goes*". This paper has adopted a different approach, one that advances a novel, parsimonious and testable framework that explains the consensus underreaction, individual overreaction, and new patterns in the data.

This original conceptualization of information offers three main advantages: (*i*) it allows for formal statistical testing of the model's restrictions, key to validate its mechanisms, especially when scrutinizing framework-specific implications (e.g., eq. (23)); (*ii*) it achieves a clear separation between information and strategic considerations, enhancing interpretability; and (*iii*) it abstains from the ubiquitous assumption of Gaussian distributed signals, which I show are empirically implausible and yet very relevant in deriving and testing implications of commonly used frameworks (Rigos (2022)).

Agents aiming to *stand out* from the crowd rationally overweigh their private information to differentiate themselves, consistent with the observed data. The analysis reveals non-negligible values for the weight of strategy in forecasters' preferences (ϕ), averaging between -0.4 and -0.7 across variables, horizons and degrees of information heterogeneity. These findings underscore the substantial role strategy plays in forecasting and caution against using forecasts as direct measures of expectations in tests of FIRE.

State-of-the-art statistical testing, adapting the methodology from Hansen (1982), supports the model's interpretation of both existing and newly presented evidence, validating the imposed restrictions crucial for assessing the economic mechanism. Importantly, the test corroborates implications that are *unique* to the proposed information structure, further validating its suitability and distancing it from the common noisy information framework. Overall, this confirms the need of strategic considerations in explaining the observed patterns.

In summary, this study highlights that assuming independence across information sets and motives of forecasters is overly restrictive and implausible in professional forecasting. Survey analysts and policymakers need to account for strategic behavior when interpreting forecasts, exercising prudence in using surveys as unbiased sources of expectations.

References

Adam, K., Kuang, P., & Xie, S. (2024). *Overconfidence in private information explains biases in professional forecasts* (tech. rep.). Mannheim University mimeo.

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

Angeletos, G.-M., Huo, Z., & Sastry, K. A. (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual*, *35*(1), 1–86.

Angeletos, G.-M., & Lian, C. (2016). Incomplete information in macroeconomics: Accommodating frictions in coordination. In *Handbook of macroeconomics* (pp. 1065–1240, Vol. 2). Elsevier.

Angrist, J. D., & Newey, W. K. (1991). Over-identification tests in earnings functions with fixed effects. *Journal of Business & Economic Statistics*, *9*(3), 317–323.

Barberis, N., Greenwood, R., Jin, L., & Shleifer, A. (2018). Extrapolation and bubbles. *Journal of Financial Economics*, *129*(2), 203–227.

Benjamini, Y., Krieger, A. M., & Yekutieli, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, *93*(3), 491–507.

Bernhardt, D., Campello, M., & Kutsoati, E. (2006). Who herds? *Journal of Financial Economics*, *80*(3), 657–675.

Bianchi, F., Ludvigson, S. C., & Ma, S. (2022). Belief distortions and macroeconomic fluctuations. *American Economic Review*, *112*(7), 2269–2315.

Bordalo, P., Gennaioli, N., Ma, Y., & Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, *110*(9), 2748–82.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2018). Diagnostic expectations and credit cycles. *The Journal of Finance*, *73*(1), 199–227.

Broer, T., & Kohlhas, A. N. (2022). Forecaster (mis-) behavior. *Review of Economics and Statistics*, 1–45.

Carvalho, C., & Nechio, F. (2014). Do people understand monetary policy? *Journal of Monetary Economics*, *66*, 108–123.

Cavallo, A., Cruces, G., & Perez-Truglia, R. (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, *9*(3), 1–35.

Chahrour, R., & Jurado, K. (2018). News or noise? the missing link. *American Economic Review*, *108*(7), 1702–1736.

Chamberlain, G. (1984). Panel data. *Handbook of econometrics*, *2*, 1247–1318.

Chui, C. K., Chen, G., et al. (2017). *Kalman filtering*. Springer.

Clement, M. B., & Tse, S. Y. (2005). Financial analyst characteristics and herding behavior in forecasting. *The Journal of finance*, *60*(1), 307–341.

Coibion, O., & Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, *120*(1), 116–159.

Coibion, O., & 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., Stark, T., et al. (2019). Fifty years of the survey of professional forecasters. *Economic Insights*, *4*(4), 1–11.

Da Silveira, R. A., Sung, Y., & Woodford, M. (2020). *Optimally imprecise memory and biased forecasts* (tech. rep.). National Bureau of Economic Research.

D'agostino, R., & Pearson, E. S. (1973). Tests for departure from normality. empirical results for the distributions of b 2 and b. *Biometrika*, *60*(3), 613–622.

Doornik, J. A., & Hansen, H. (2008). An omnibus test for univariate and multivariate normality. *Oxford bulletin of economics and statistics*, *70*, 927–939.

ECB. (2014). Results of the second special questionnaire for participants in the ecb survey of professional forecasters. *Monthly Bulletin*, 1–28.

Eva, K., & Winkler, F. (2023). A comprehensive empirical evaluation of biases in expectation formation.

Farmer, L., Nakamura, E., & Steinsson, J. (2021). Learning about the long run.

Fuster, A., Hebert, B., & Laibson, D. (2012). Natural expectations, macroeconomic dynamics, and asset pricing. In *Nber macroeconomics annual* (Vol. 26). University of Chicago Press.

Fuster, A., Laibson, D., & Mendel, B. (2010). Natural expectations and macroeconomic fluctuations. *Journal of Economic Perspectives*, *24*(4), 67–84.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, *129*(4), 1661–1710.

Gallo, G. M., Granger, C. W., & Jeon, Y. (2002). Copycats and common swings: The impact of the use of forecasts in information sets. *IMF staff Papers*, *49*(1), 4–21.

Gemmi, L., & Valchev, R. (2023). *Biased surveys* (tech. rep.). National Bureau of Economic Research, WP 31607.

Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American economic review*, *70*(3), 393–408.

Hamilton, J. D. (1994). *Time series analysis*. Princeton University Press.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the econometric society*, 1029–1054.

Hong, H., & Kacperczyk, M. (2010). Competition and bias. *The Quarterly Journal of Economics*, *125*(4), 1683–1725.

Juodis, A., & Kučinskas, S. (2023). Quantifying noise in survey expectations. *Quantitative Economics*, *14*(2), 609–650.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology*, *3*(3), 430–454.

Kohlhas, A. N., & Walther, A. (2021). Asymmetric attention. *American Economic Review*, *111*(9), 2879–2925.

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

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

Malmendier, U., & Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, *131*(1), 53–87.

Mankiw, N. G., & Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, *117*(4), 1295–1328.

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

Milgrom, P. R., & Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, 1089–1122.

Morris, S., & Shin, H. S. (2002). Social value of public information. *american economic review*, *92*(5), 1521–1534.

Ottaviani, M., & Sørensen, P. N. (2006a). Reputational cheap talk. *The Rand journal of economics*, *37*(1), 155–175.

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

Paul, T. (2023). Measuring information frictions with surveys of expectations: How revised data bias estimates.

Prescott, E. C. (1977). Should control theory be used for economic stabilization? *Carnegie-Rochester Conference Series on Public Policy*, *7*, 13–38.

Rigos, A. (2022). The normality assumption in coordination games with flexible information acquisition. *Journal of Economic Theory*, *203*, 105485.

Sargent, T. J. (1999). *The conquest of american inflation*. Princeton University Press.

Sarte, P.-D. (2014). When is sticky information more information? *Journal of Money, Credit and Banking*, *46*(7), 1345–1379.

Scheinkman, J. A., & Xiong, W. (2003). Overconfidence and speculative bubbles. *Journal of political Economy*, *111*(6), 1183–1220.

Shapiro, S. S., & Wilk, M. B. (1965). An analysis of variance test for normality (complete samples). *Biometrika*, *52*(3-4), 591–611.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica: journal of the Econometric Society*, 1–48.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, *50*(3), 665–690.

Woodford, M. (2001). Imperfect common knowledge and the effects of monetary policy.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

Appendices

A Variable construction

The construction of each variable follows the established routine by Bordalo et al. (2020). For each variable, I report the survey time, the survey question, and the definitions of forecast variable (year-over-year growth, following the literature), revision variable, and actual variable.

1. NGDP

Variable	Nominal GDP. Source: SPF.
Time	Around the 3rd week of the middle month in the quarter.
Question	The level of nominal GDP in the current quarter and the next 4 quarters.
Forecast	Nominal GDP growth from end of quarter $t-1$ to end of quarter $t+h$: $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1$, where t is the quarter of forecast and x is the level of GDP in a given quarter; x_{t-1} uses the initial release of actual value in quarter $t-1$, available by the time of the forecast in quarter t.
Revision	$\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}$.
Actual	$\frac{x_{t+h}}{x_{t-1}} - 1$, using real-time macro data published in quarter $t+4$.

2. RGDP

Variable	Real GDP. Source: SPF.
Time	Around the 3rd week of the middle month in the quarter.
Question	The level of real GDP in the current quarter and the next 4 quarters.
Forecast	Real GDP growth from end of quarter $t-1$ to end of quarter $t+h$: $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1$, where x_{t-1} is the initial release of quarter $t-1$.
Revision	$\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}$.
Actual	$\frac{x_{t+h}}{x_{t-1}} - 1$, using real-time macro data published in quarter $t+4$.

3. PGDP

Variable	GDP price deflator. Source: SPF.
Time	Around the 3rd week of the middle month in the quarter.
Question	The level of GDP price deflator in the current quarter and the next 4 quarters.
Forecast	GDP deflator inflation from end of quarter $t-1$ to $t+h$: $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1$, with x_{t-1} the initial release of quarter t.
Revision	$\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}$.
Actual	$\frac{x_{t+h}}{x_{t-1}} - 1$, using real-time macro data published in quarter $t+4$.

4. CPI

Variable	Consumer Price Index. Source: SPF.
Time	Around the 3rd week of the middle month in the quarter.
Question	CPI growth rate in the current quarter and the next 4 quarters.
Forecast	$\mathbb{F}_t\Big[\prod_{j=0}^{3}(z_{t+j}/4+1)\Big]$; simple average $\mathbb{F}_t\,\frac{z_t+z_{t+1}+z_{t+2}+z_{t+h}}{4}$ similar.
Revision	Difference of successive forecasts of that product.
Actual	$\frac{x_{t+h}}{x_{t-1}} - 1$, real-time data (pre-1994Q3 uses 1994Q3 vintage).

5. RCONSUM

Variable Real consumption. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The level of real consumption in the current quarter and the next 4 quarters.

Forecast $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1.$

Revision $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}.$

Actual $\frac{x_{t+h}}{x_{t-1}} - 1$, real-time data published in quarter $t+4$.

6. INDPROD

Variable Industrial production index. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question Average level of the industrial production index in the current quarter and the next 4 quarters.

Forecast $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1.$

Revision $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}.$

Actual $\frac{x_{t+h}}{x_{t-1}} - 1$, real-time data published in quarter $t+4$.

7. RNRESIN

Variable Real non-residential investment. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The level of real non-residential investment in the current quarter and the next 4 quarters.

Forecast $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1.$

Revision $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}.$

Actual $\frac{x_{t+h}}{x_{t-1}} - 1$, real-time data published in quarter $t+4$.

8. RRESIN

Variable Real residential investment. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The level of real residential investment in the current quarter and the next 4 quarters.

Forecast $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1.$

Revision $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}.$

Actual $\frac{x_{t+h}}{x_{t-1}} - 1$, real-time data published in quarter $t+4$.

9. RGF

Variable Real federal government consumption. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The level of real federal government consumption in the current quarter and the next 4 quarters.

Forecast $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1.$

Revision $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}.$

Actual $\frac{x_{t+h}}{x_{t-1}} - 1$, initial x_{t+h} released in quarter $t+4$.

10. RGSL

Variable Real state & local government consumption. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The level of real state & local government consumption in the current quarter and the next 4 quarters.

Forecast $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1.$

Revision $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}.$

Actual $\frac{x_{t+h}}{x_{t-1}} - 1$, real-time data published in quarter $t+4$.

11. HOUSING

Variable Housing starts. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The level of housing starts in the current quarter and the next 4 quarters.

Forecast $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - 1.$

Revision $\frac{\mathbb{F}_t\, x_{t+h}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+h}}{\mathbb{F}_{t-1}\, x_{t-1}}.$

Actual $\frac{x_{t+h}}{x_{t-1}} - 1$, real-time data published in quarter $t+4$.

12. UNEMP

Variable Unemployment rate. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The average unemployment rate in the current quarter and the next 4 quarters.

Forecast $\mathbb{F}_t\, x_{t+h}$.

Revision $\mathbb{F}_t\, x_{t+h} - \mathbb{F}_{t-1}\, x_{t+h}$.

Actual x_{t+h}, real-time data published in quarter $t+4$.

13. TB3M

Variable 3-month Treasury rate. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The average 3-month Treasury rate in the current quarter and the next 4 quarters.

Forecast $\mathbb{F}_t\, x_{t+h}$.

Revision $\mathbb{F}_t\, x_{t+h} - \mathbb{F}_{t-1}\, x_{t+h}$.

Actual x_{t+h}.

14. TN10Y

Variable 10-year Treasury rate. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The average 10-year Treasury rate in the current quarter and the next 4 quarters.

Forecast $\mathbb{F}_t\, x_{t+h}$.

Revision $\mathbb{F}_t\, x_{t+h} - \mathbb{F}_{t-1}\, x_{t+h}$.

Actual x_{t+h}.

15. AAA

Variable AAA corporate bond rate. Source: SPF.

Time Around the 3rd week of the middle month in the quarter.

Question The average AAA corporate bond rate in the current quarter and the next 4 quarters.

Forecast $\mathbb{F}_t\, x_{t+h}$.

Revision $\mathbb{F}_t\, x_{t+h} - \mathbb{F}_{t-1}\, x_{t+h}$.

Actual x_{t+h}.

A.1 Real-Time Methodology

First, the real-time methodology is indeed a relevant aspect in the evaluation of forecasting, and the main goal of researcher is to align the data with the forecasters' information sets at the time of forecast issuance. As described Appendix A, I focus on *initial* releases sourced from the Philadelphia Fed's Real-Time Dataset for Macroeconomists. Notice that this is *not* implying that I use initial estimates to compute the actual realizations, but only to align each forecaster's beliefs with the most recent information. Then:

- Revision: $\frac{\mathbb{F}_t\, x_{t+3}}{x_{t-1}} - \frac{\mathbb{F}_{t-1}\, x_{t+3}}{\mathbb{F}_{t-1}\, x_{t-1}}$.
- Actual: $\frac{x_{t+3}}{x_{t-1}} - 1$, using real time macro data published in quarter $t + 4$.

with x_{t-1} being initial release in quarter $t - 1$, available by the time of the forecast in quarter t.

A.2 Early Sample

There exist concerns about the reliability of the first years (1968-1971) of the SPF. To ensure robustness of the patterns displayed in Section 2, I rerun the regressions excluding the early years (before 1971). Figure 10 shows some of these results:

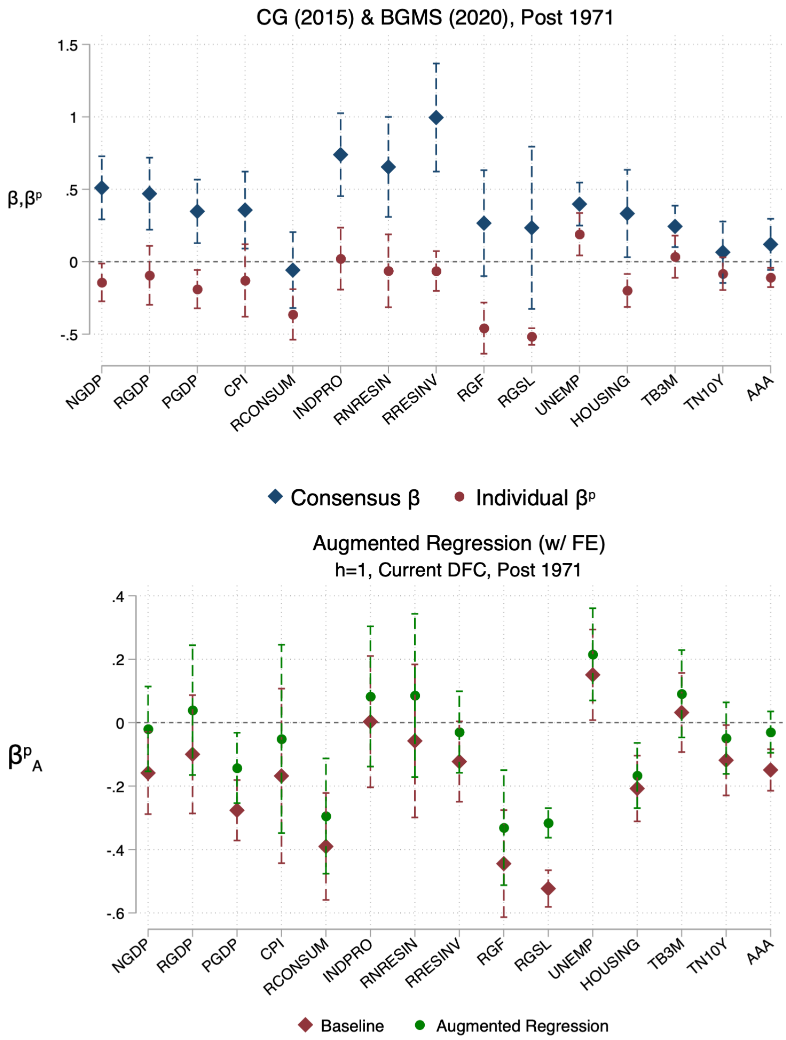

Figure 10: Consensus, individual and augmented regressions: post 1971 sample.

B Strategic Forecasting: Literature Scheme

B.1 Strategic Complementarity

Olsen (1996), *FAJ*

Setting: Professional finance forecasts (EPS forecasts)
Idea: Behavioral desire to conform
Facts: $\text{Var}(\mathbb{F}(x)) < \text{Var}(x)$; $\mathbb{E}(\mathbb{F}(x) > \mathbb{E}(x)$
Result: Forecasters herd. Their forecasts match a behavioral explanation.

Gallo et al. (2002), *IMF Working Papers*

Setting: Professional macro forecasts
Idea: Copycat behavior
Facts: $\text{Var}(\mathbb{F}(x))$ is a negative function of the forecast horizon
Result: Forecasters herd.

Krishnam et al. (2006), *AAA*

Setting: Professional finance forecasts
Idea: Directly estimate true posterior and observe deviation. Model: forecasters know prior consensus and bias their forecast away from objective expectation (posterior).
Facts: NA
Result: Forecasters herd.

Hong and Kacperczyk (2010), *QJE*

Setting: Professional finance forecasts
Idea: Merge of brokerage houses produces excess forecasters. Study competition effects among them.
Facts: Forecasts improve when there is *skin in the game* (less bias)
Result: Forecasters herd more when career concerns are present.

Hong and Kubik (2003), *JoF*

Setting: Professional finance forecasts
Idea: Agency problem between employer and forecaster: different agenda.
Facts: Negative correlation between forecasts' accuracy and job separation. More so when forecasters are juniors. Positive correlation between forecast boldness and job separation. More so when forecasters are juniors. These define "implicit incentives".
Result: Forecasters herd when career concerns are present.

Clement and Tse (2005), *JoF*

Setting: Professional finance forecasts
Idea: Extend Hong and Kubik (2003) to more observables.
Facts: Forecasts' boldness is a increasing function of: (i) seniority; (ii) brokerage house size; (iii) prior accuracy; (iv) forecast horizon. All of these matter to the forecaster as career incentives.
Result: Forecasters herd when career concerns are present.

Lamont (2002), ***JEB&O***

Setting: Professional macro forecasts
Idea: Agency problem between employer and forecaster: different agenda. Reputation concerns.
Facts: Negative correlation between forecasts' seniority and herding.
Result: Forecasters herd when career concerns are present.

Jagadeesh and Kim (2010), ***RFS***

Setting: Professional finance forecasts.
Idea: Model compensation as a function of accuracy.
Facts: Positive correlation between stock price and revisions. (Not clear what's the link)
Result: Forecasters herd.

B.2 Strategic Substitutability

Bernhardt et al. (2006), ***JFE***

Setting: Professional finance forecasts.
Idea: Statistical test. If there is herding, more than 50% of the time the forecast will fall between the consensus and the realization.
Facts: The opposite is true. Individual forecasts are less likely to be between consensus and realization
Result: Forecasters anti-herd.

Ottaviani and Sorensen (2006), ***JFE***

Setting: Professional finance forecasts.
Idea: (i) Reputational cheap talk. It posits that forecasters endeavor to convince the market that they are well informed. (ii) Winner takes all contest.
Facts: NA
Result: Forecasters anti-herd.

Ashiya (2009), ***Journal of Forecasting***

Setting: Japanese professional macro forecasts.
Idea: Forecasters' wages are based on accuracy and publicity, so they will make more extreme and less accurate forecasts.
Facts: Industry-based distinction between groups (trading, research, banking...)
Result: Forecasters anti-herd.

Laster et al. (1999), ***QJE***

Setting: Professional macro forecasts.
Idea: Forecasters' wages are based on accuracy and publicity, so they will make more extreme and less accurate forecasts.
Facts: Industry-based distinction between groups (trading, research, banking...)
Result: Forecasters anti-herd.

B.3 Ambiguous Strategic Incentives

Ehrbeck and Waldmann(1996), *QJE*

Setting: Professional finance forecasts.
Idea: Forecasters rationally bias their forecasts strategically if they systematically bias them toward the best forecasters in the market. Otherwise, it's behavioral.
Facts: NA
Result: Forecasters are behaviorally biased

B.4 Deviation from min{MSE} in equilibrium

One classic critique to the survival of strategic incentives in equilibrium can be summarized as follows:

> *"Forecasters have customers who pay for their forecasts. Shouldn't these customers pay attention to the performances of forecasters and therefore exert discipline on their loss function?"*

Essentially, this reduces to the characterization of the equilibrium in the presence of a rational expectations equilibrium (REE) solution concept. Even if there is no existing evidence on how the markets for forecasts work (due to the absence of data on client-forecaster relationship), these paragraphs provide a simple heuristic that supports the possibility of strategic incentives in equilibrium. The key intuition relies in the fact that, in the professional forecasting game, *accuracy* does not pay a high reward when it is diluted among every player. That is, there exist a trade-off between accuracy and *publicity* such that "being right" whenever everyone else is is a low return strategy, leading forecasters to pursue riskier forecasting strategies (especially when interested in expanding their customer base).

One might argue that as long as there exists symmetry in the payoff function of professional forecasters this mechanism should still converge toward a *truthful revelation* equilibrium. That is, that as long as "being right & different" is as rewarding as "being wrong & different" is penalizing, forecasters' optimal strategy should be centered around providing their best estimate of the actual realization, i.e. announcing their true expectation. A graphical heuristic follows: Then, an asymmetric return function would deliver

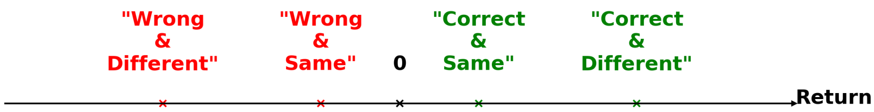

Figure 11: Symmetric return function

the strategic incentive to deviate from truthful revelation in order to *stand out* from the crowd: The

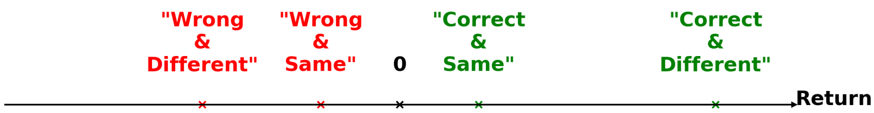

Figure 12: Asymmetric return function

anecdotal evidence I collected through interviews is supportive of the second (asymmetric) case. A large, formal literature in game theory (surveyed by Marinovic et al. (2013)) argues around these types of incentives. They condense the central point in the following excerpt:

> *Then, at* $\mathbb{E}(x|s)$*, it is optimal to deviate to issuing a forecast which is closer to* s [the private signal, Ed.] *than* $\mathbb{E}(x|s)$ *because the first-order reduction in the expected number of winners*

with whom the prize must be shared more than compensates the second-order reduction in the probability of winning.

B.4.1 Nash Equilibrium under Loss Function (8)

Let $N \geq 2$ agents announce forecasts $\mathbb{F}(x_{t+h}|\Omega_t^\tau) \equiv f_\tau$. Write the cross–sectional average $\bar{\mathbb{F}}_t(x_{t+h}) = \frac{1}{N}\sum_{\iota=1}^{N} f_\iota$ and $S_\tau \equiv \sum_{\iota\neq\tau} f_\iota$. Agent τ minimizes

$$\mathcal{L}_\tau(f_\tau, f_{-\tau}) = \mathbb{E}\Big[\big(x_{t+h} - f_\tau\big)^2 + \phi\big(\bar{\mathbb{F}}_t(x_{t+h}) - f_\tau\big)^2 \,\Big|\, \Omega_t^\tau\Big]. \tag{33}$$

Since $\bar{\mathbb{F}}_t = \dfrac{f_\tau + S_\tau}{N}$, set $\alpha \equiv \dfrac{N-1}{N}$ and denote $\mu_\tau \equiv \mathbb{E}(x_{t+h}|\Omega_t^\tau)$. The conditional first derivative is

$$\frac{\partial \mathcal{L}_\tau}{\partial f_\tau} = -2\big(\mu_\tau - f_\tau\big) - 2\phi\frac{N-1}{N}\Big[\alpha f_\tau - \tfrac{S_\tau}{N}\Big].$$

Setting it to zero and solving gives the best response

$$f_\tau^\star = \frac{\mu_\tau + \phi\alpha^2 \dfrac{S_\tau}{N-1}}{1+\phi\alpha^2}. \tag{34}$$

Agent τ observes S_τ only through the conditional expectation of the consensus: $\mathbb{E}_\tau\big(\bar{\mathbb{F}}_t(x_{t+h})\big) = \dfrac{f_\tau + S_\tau}{N}$. Eliminating S_τ in (34) yields

$$\mathbb{F}^*(x_{t+h}|\Omega_t^\tau) = \frac{1}{1+\phi\big(1-\frac{1}{N}\big)^2}\,\mathbb{E}(x_{t+h}|\Omega_t^\tau) + \frac{\phi\big(1-\frac{1}{N}\big)}{1+\phi\big(1-\frac{1}{N}\big)^2}\,\mathbb{E}_\tau\big(\bar{\mathbb{F}}_t(x_{t+h})\big) \tag{35}$$

where the differences from the baseline (9) are driven by the fact that a rigorous characterization of optimality considers the effect of each forecaster's optimal forecast on the consensus forecast[42]

The second-order condition amounts to

$$\frac{\partial^2 \mathcal{L}_\tau}{\partial f_\tau^2} = 2 + 2\phi\frac{(N-1)^2}{N^2} > 0 \quad \forall\phi > -1,$$

so (35) minimises (33).

Symmetric Nash equilibrium. Suppose every agent follows (35). Because each $f_\tau^\star$ is a best response to $f_{-\tau}$ and the second-order test is satisfied, the profile $\big\{\mathbb{F}^*(x_{t+h}|\Omega_t^\tau)\big\}_{\tau=1}^{N}$ constitutes a Nash equilibrium.

B.5 *Cheap talk*: Anonymity and Surveys

Anonymity poses a threat to surveys by uncoupling survey responses and respondents. This issue falls within the realm of the so-called *cheap talk* critique, the belief that costless, non-binding communication is ineffective and/or non-credible in accurately conveying private information and influencing decisions in strategic settings. The SPF is an anonymous survey, and yet studies like Lamont (2002) and Laster

[42] I explore this alternative scenario numerically and it makes little to no difference (estimated ϕ is slightly smaller across variables). Therefore, I treat (9) as baseline for its ease of interpretation.

et al. (1999) argued in favor of a strategic bias in the responses of survey participants. Croushore, Stark, et al. (2019) responded: "We never publish a panelist's name with their projection. In principle, this policy removes the potential for a publicity motive affecting the projections". This paper challenges the efficacy of anonymity in preserving truthful responses in professional forecasts.

In fact, Marinovic et al. (2013) directly addressed the debate by claiming that respondents treat anonymous and non-anonymous surveys identically. Their argument rely on two key observations: (*i*) forecast production is a costly activity, constituting a disincentive for forecasters to specialize responses for anonymous and non-anonymous surveys; (*ii*) there is qualitative evidence that forecasters are concerned about "strategy detection" by the editors of the surveys they participate in, detection that would inevitably lead to their exclusion from the survey. Furthermore, BGMS are aware of the potential concern anonymity implies and use the Blue Chip Economic Indicators, a non-anonymous survey, as complementary confirmation of their empirical findings. Indeed, the evidence they present is almost indistinguishable across Blue Chip and SPF for the majority of variables. This is in line with Marinovic et al. (2013)'s hypotheses. Additionally, Gemmi and Valchev (2023) provide further corroboration using the Federal Reserve's *Tealbooks*, showing that strategy cannot be detected in non-professional forecasting. Finally, the European Central Bank has tackled the issue explicitly in a recent study (ECB (2014)) concerning their own Survey of Professional Forecasters. They ask respondents: *"When responding to the SPF, what forecast do you provide?"*. 84% of respondents chose the option *"provide latest available forecast"*. The same questionnaire asks: *"Do you publish externally the forecasts that you send to the ECB?"*[43]. 86% of the interviewed replied affirmatively.

Lastly, let me point out that the identifiability of professional forecasters' responses is not *necessitated* to generate strategic incentives. In fact, most of the studies occupied with providing the micro-foundations of strategic motives[44] only rely on the *internal* (within company) observability of the forecasting performance. In other words, the so-called "implicit incentive schemes" work through principal-agent problems that happen within the single firm or forecasting team: the only necessary premise is the presence of divergent incentives between the person/group that formulates the announced forecast and their employer. That alone is sufficient to create the reputational/career concerns generating strategy motive.

[43] Questions 1b and 14 of the questionnaire.

[44] Refer to the previous section and to Appendix B for a comprehensive overview.

C Individual β under rationality

The following proof shows that, under the assumption of individual rationality (or *Bayesianism*) *and* accuracy-maximizer forecasters ($\mathbb{E} \equiv \mathbb{F}$), the coefficient on forecast revisions is always zero, regardless of the noise structure of signals. The proof is simplified to the basic static case, but is without loss of generality.

$$\theta \sim N\left(\mu_0, \sigma_\theta^2\right)$$
$$s = \theta + \varepsilon \quad \varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$$

$$\begin{aligned} FE &= \theta - \mathbb{E}(\theta \mid s) = \theta - \mathbb{E}(\theta) - K(s - \mathbb{E}(\theta)) \\ &= \theta - \mathbb{E}(\theta) - K(\theta + \varepsilon - \mathbb{E}(\theta)) \\ &= \theta(1-K) - K\varepsilon - \mathbb{E}(\theta)(1-K) \end{aligned}$$

$$\begin{aligned} FR &= \mathbb{E}(\theta \mid s) - \mathbb{E}(\theta) = \mathbb{E}(\theta) + K(s - \mathbb{E}(\theta)) - \mathbb{E}(\theta) \\ &= K(\theta + \varepsilon - \mathbb{E}(\theta)) \end{aligned}$$

$$\begin{aligned} FE &= \alpha + \beta FR + e \\ \beta &= \frac{\mathrm{Cov}(FE, FR)}{\mathrm{Var}(FR)} = \frac{\mathrm{Cov}\left[\theta(1-K) - K\varepsilon - \mathbb{E}(\theta)(1-K), K(\theta + \varepsilon - \mathbb{E}(\theta))\right]}{\mathrm{Var}\left[K(\theta + \varepsilon - \mathbb{E}(\theta))\right]} \\ &= \frac{(1-K)K\,\mathrm{Cov}(\theta,\theta) - K^2\,\mathrm{Cov}(\varepsilon,\varepsilon) - (1-K)K\,\mathrm{Cov}(\mathbb{E}(\theta),\mathbb{E}(\theta))}{K^2\left(\sigma_\theta^2 + \sigma_\varepsilon^2\right)} \\ &= \frac{(1-K)\sigma_\theta^2 - K\sigma_\varepsilon^2}{K\left(\sigma_\theta^2 + \sigma_\varepsilon^2\right)} \qquad \text{now use } K = \frac{\sigma_\varepsilon^{-2}}{\sigma_\theta^{-2} + \sigma_\varepsilon^{-2}} \\ &= \frac{\frac{\sigma_\theta^{-2} + \sigma_\varepsilon^{-2} - \sigma_\varepsilon^{-2}}{\sigma_\theta^{-2} + \sigma_\varepsilon^{-2}} \cdot \sigma_\theta^2 - \frac{\sigma_\varepsilon^{-2}}{\sigma_\theta^{-2} + \sigma_\varepsilon^{-2}}\sigma_\varepsilon^2}{\frac{\sigma_\varepsilon^{-2}}{\sigma_\theta^{-2} + \sigma_\varepsilon^{-2}}\left(\sigma_\theta^2 + \sigma_\varepsilon^2\right)} \\ \beta &= \frac{1}{\sigma_\theta^2}\left(\frac{1}{\sigma_\theta^{-2} + \sigma_\varepsilon^{-2}} - \frac{1}{\sigma_\theta^{-2} + \sigma_\varepsilon^{-2}}\right) = 0 \end{aligned}$$

Another way to read this result is the following: while $K = 1 \implies \beta = 0$ (this holding also at the consensus level), the opposite is not true: $\beta = 0 \not\Longrightarrow K = 1$. In other words, the null hypothesis of rationality is uninformative about the quality of information in noisy models, and only accounts for the optimization routine analysts employ when maximizing accuracy.

D Augmented Regression: a closer look

D.1 Timing of consensus' observation

A central tenet in the context of strategic forecasting is that forecasters consider their positioning relative to the contemporaneous consensus response (equal-weighted average of individual responses) when submitting their estimates. An intuitive critique concerns the timing of the observability the consensus estimate. More precisely, forecasters might observe the consensus once they posted their forecast and the survey results are released, so that they cannot directly infer the contemporaneous deviation of their forecasts from the consensus.

I will articulate my answer to this issue in two parts: on the evidence supporting the observability of *contemporaneous* consensus responses; and on the substantial irrelevance of such point, both theoretically and empirically. Moreover, regardless of the cogency of the first two arguments, I will include additional empirical solutions that obviate to the problem in the first place, like using a "lagged" and a "mixed" measure of DFC.

First, let me notice that the ideal (unfeasible) regression we would like to run is the following:

$$\underbrace{x_{t+h} - \mathbb{F}_t^i(x_{t+h})}_{\text{Forecast Error}} = \alpha_A^{p,i} + \beta_A^p \Big(\underbrace{\mathbb{F}_t^i(x_{t+h}) - \mathbb{F}_{t-1}^i(x_{t+h})}_{\text{Forecast Revision}} \Big) + \gamma_A^p \Big(\underbrace{\mathbb{F}_t^i(x_{t+h}) - \mathbb{E}^i \left(\overline{\mathbb{F}}_t(x_{t+h}) \right)}_{\textit{Expected} \text{ Deviation from Consensus}} \Big) + u_{A,t}^{p,i} \quad (36)$$

Being (36) de facto unobservable (i.e. the unobservability of expectations is indeed the core premise of the paper), the task is to argue around the best way to capture the core mechanisms behind the exercise. In fact, the idea is that forecasters react to the contemporaneous cross-sectional average, not the past one. If assuming that SPF participants are aware of the consensus' stance is realistic, than that should be the preferable strategy to any proxy. I will argue that a close approximation to the consensus is in the information set of forecasters, because (i) professional forecasters simultaneously participate to multiple surveys that are not synchronized, and therefore can observe the consensus estimates in other pools before submitting their response. Other surveys serve as an almost perfect indicator of the current consensus; (ii) a model with a large number of types would imply that the consensus forecast evolves only as a function of the current realization x_t and the public signal p_t, as idiosyncratic information washes out on average. The same holds true for a model where private information is nonexistent or very noisy; (iii) finally, consensus forecasts are very persistent processes, a fact that combined with the observation of correlated information at t might suggest that t's averages are easily estimated by forecasters.

Secondly, on the likely irrelevance of the debate: conceptually, the augmented regression is employed to argue that the estimates of β^p in equation (Body-2) are biased, and therefore cannot be mapped into measures of information rigidity nor behavioral biases. Reiterating from the body of the paper, if the error term in equation (Body-2) is orthogonal to the forecast revision, introducing any other variable to the regression model should not systematically alter the estimation of β_A^p. On the contrary, across all variables, the point estimates of β_A^p in the augmented regression are consistently higher than in the baseline specification. As I point out below, the bias holds true regardless of the use of current or past DFC (although, the sign flips, and I explain the reason below). The general point is that DFC consistently shift the estimation of the β_A^p, suggesting a misspecification in the standard test of FIRE.

Finally, below I display the results using alternative feasible measures of DFC:

$$\underbrace{x_{t+h} - \mathbb{F}_t^i(x_{t+h})}_{\text{Forecast Error}} = \alpha_A^p + \beta_A^p \Big[\underbrace{\mathbb{F}_t^i(x_{t+h}) - \mathbb{F}_{t-1}^i(x_{t+h})}_{\text{Forecast Revision}} \Big] + \gamma_A^p \Big[\underbrace{\mathbb{F}_{t-1}^i(x_{t+h}) - \overline{\mathbb{F}}_{t-1}(x_{t+h})}_{\text{Deviation f. Consensus}} \Big] + u_t^{p,i} \quad (37)$$

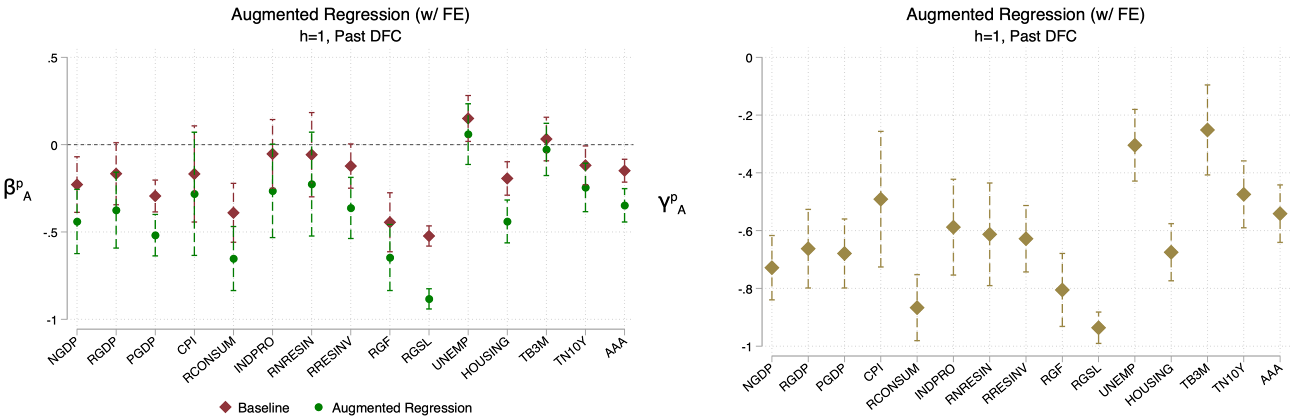

Figure 13: Lagged-DFC version of the augmented regression.

As it is immediate to notice, the γ^p is still negative and significantly different from zero, while for the β_A^p the bias is still there, with the sign flipped: the baseline is consistently higher than the augmented estimation. Why is that? Well, it is easy to show that the sign of the bias is a function of the coefficient on the omitted regressor (the γ_A^p, in this case) and the covariance between the independent variable and the omitted regressor ($\text{Cov}(FR, DFC)$):

$$\mathbb{E}(\hat{\beta}^p) \xrightarrow{p} \beta^p + \gamma^p \frac{\text{Cov}(FR, DFC)}{\text{Var}(FR)} \tag{38}$$

When using *current* DFC, the covariance term is positive because $\mathbb{F}_t^i(x_{t+h})$ appears with the same sign in FR ($\mathbb{F}_t^i(x_{t+h}) - \mathbb{F}_{t-1}^i(x_{t+h})$) and in DFC ($\mathbb{F}_t^i(x_{t+h}) - \overline{\mathbb{F}}_t(x_{t+h})$). Combined with $\hat{\gamma}^p < 0$, we conclude that there is a *downward* bias. The opposite becomes true if the use the *past* DFC ($\mathbb{F}_{t-1}^i(x_{t+h}) - \overline{\mathbb{F}}_{t-1}(x_{t+h})$), as now the covariance will be negative.

Further confirmation comes from an exercise using "mixed" distance from consensus: $\mathbb{F}_t^i(x_{t+h}) - \overline{\mathbb{F}}_{t-1}(x_{t+h})$:

$$\underbrace{x_{t+h} - \mathbb{F}_t^i(x_{t+h})}_{\text{Forecast Error}} = \alpha_A^p + \beta_A^p \Big[\underbrace{\mathbb{F}_t^i(x_{t+h}) - \mathbb{F}_{t-1}^i(x_{t+h})}_{\text{Forecast Revision}}\Big] + \gamma_A^p \Big[\underbrace{\mathbb{F}_t^i(x_{t+h}) - \overline{\mathbb{F}}_{t-1}(x_{t+h})}_{\text{Mixed DFC}}\Big] + u_t^{p,i} \tag{39}$$

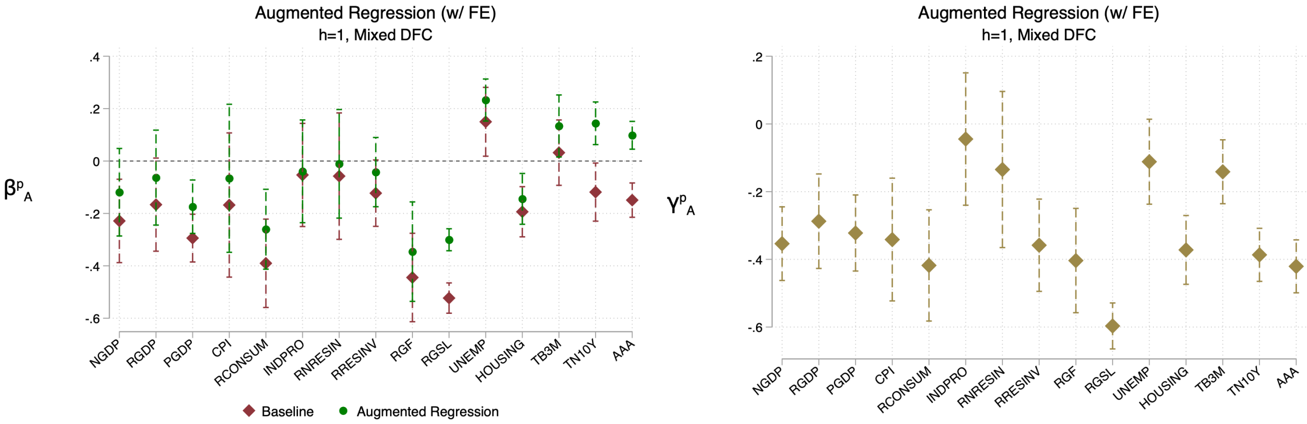

Figure 14: Lagged-consensus version of the augmented regression ("mixed DFC").

Lastly, let me note that the approach taken in the literature on related issues (Bordalo et al. (2020), Kohlhas and Walther (2021), Gemmi and Valchev (2023)) has been to proxy information known to the public with past consensus estimates[45], which I argue is a dominated strategy in light of the fact that it ignores the information flow between $t-1$ and t that updates forecasters' information sets. Such approach would "miss" on any recent development able to shift the stance of the consensus. Even if it could not be observed *directly* – a point I have attempted to refute – I would argue that the actual current consensus estimate is a more reliable measure of forecasters' beliefs on the average response.

D.2 Asymmetric DFC

A case could be made for the relevance of an asymmetric effect of the deviation from consensus on the forecast errors. I control for such scenario employing the following specification:

$$x_{t+h} - \mathbb{F}^i_t(x_{t+h}) = \alpha^p_A + \beta^p_A FR^i_t + \gamma^p_{A,+} DFC^i_{t,+} + \gamma^p_{A,-} DFC^i_{t,-} + \epsilon^i_{t+h} \tag{40}$$

Unsubstantial changes from the baseline defer the necessity for substantial discussion on these results.

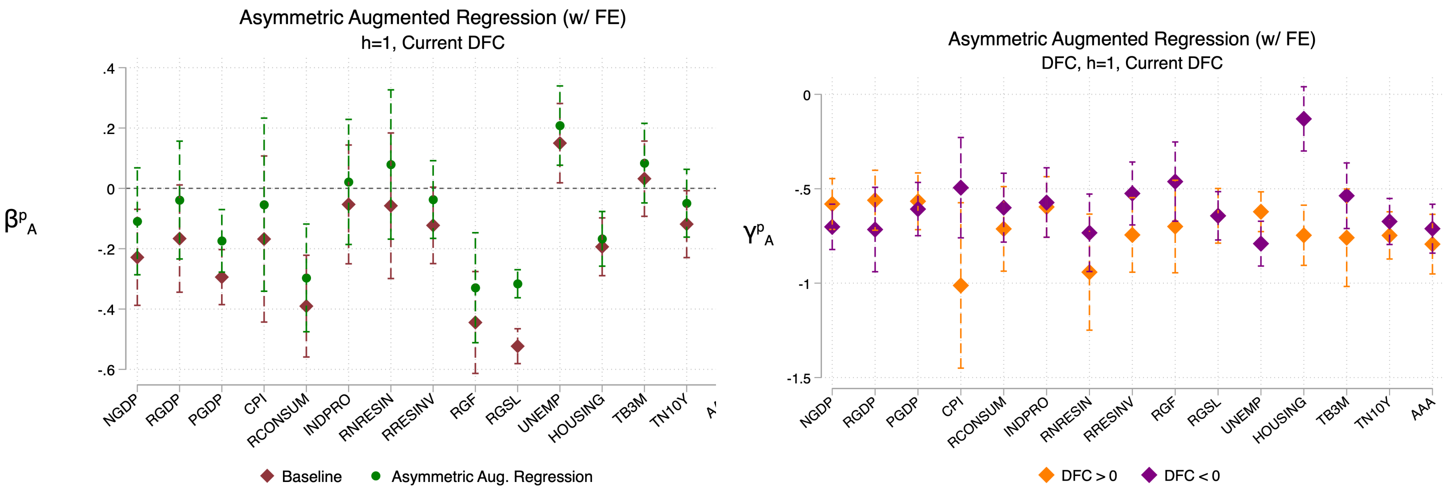

(a) β^p_A estimates for regression (40). **(b)** $\gamma^p_{A,+}, \gamma^p_{A,-}$ estimates for regression (40).

Figure 15: Regression estimates for asymmetric DFC analysis.

D.3 $\gamma^p_A \neq 0$: alternative interpretations

One could argue that $\gamma^p_A \neq 0$ might be consistent with forecasters' truthful revelation of their expectations. The model-implied expression for γ^p_A is different from zero even when $\phi = 0$ (to verify, see the expression for the structural coefficient in Appendix H). Let me point to the causes of the unintuitive result through the lens of the model. Two crucial assumptions generate it. First, there exist information *types*, and these obviously correlate with themselves. Second, types are in finite number, N. Even when $\phi = 0$ (that is, when forecasters solve `min{MSE}` problems), the coefficient on **current** deviation from consensus is negative ($\gamma^p_A < 0$). This should not come as a surprise, as even in the absence of strategic motives, the nature of information is incomplete, and this generates the inertia in beliefs that a negative coefficient

[45] Bordalo et al. (2020), p.2761: "*In Appendix C, Table C7 we address this mechanism by controlling in the pooled specification of equation (2) for the deviation of the forecast in quarter* $t-1$ *from the consensus* $(x^i_{t+h|t-1} - x_{t+h|t})$. *The consensus is released between quarter* $t-1$ *and quarter* t*, so controlling for the deviation takes into account potential news and adjustments related to the release of the consensus.*"

implies: as long as N is finite, private signals do not *wash out* completely in aggregation, and their negative covariance emerges in γ_A^p. The role of ϕ is then one of *amplifier/dampener* of the non-null coefficient.
To see this in a simpler (univariate) context, notice that the same happens in the CG regression, where β^{CG} remains positive even when $\phi = 0$. There, ϕ *amplifies* the (already positive) coefficient.

Moreover, I address such concern by implementing a slightly modified version of (Body-3) (as discussed in Appendix D.1, that is, using a "mixed" deviation from consensus measure (equation (39)). Using the lagged-consensus delivers unequivocally $\gamma_A^p = 0$. In fact, the new mapping to the model-implied expression follows:

$$\gamma_A^p = \frac{\frac{DQ}{N}\sigma_y^2\left[\sigma_y^2\left(\rho^2Q^2+\rho^2(N-1)+D^2\right)+\rho^2\sigma_\epsilon^2+\sigma_p^2\right]-\left[\sigma_p^2+\sigma_y^2\left(\Lambda Q+\Lambda(N-1)+\frac{D^2}{N}\right)\right]QD\sigma_y^2}{\left[\sigma_y^2\left(\rho^2Q^2+\rho^2(N-1)+D^2\right)+\rho^2\sigma_\epsilon^2+\sigma_p^2\right]\left[\sigma_p^2+\sigma_\epsilon^2+\sigma_y^2\left(N\Lambda+\left(\frac{D}{N}\right)^2\right)\right]-\left[\sigma_p^2+\sigma_y^2\left(\Lambda Q+\Lambda(N-1)+\frac{D^2}{N}\right)\right]^2} \tag{41}$$

with $\Lambda = \frac{N^2-N-\phi(N-1)}{N}$.

The key takeaway from (41) is that when $\phi = 0$, $\gamma_A^p = 0$, reconciling our intuition with the internal consistency of the model.[46] The empirical estimates of such specification are presented in Figure 14. It is evident from the right panel that $\hat{\gamma}_A^p < 0$ holds in the great majority of cases.

[46] What is happening "in the background" is that the forces that were operating towards a negative coefficient – byproduct of the incomplete information environment – are now non-impactful, as there exists a lag between the positive-horizon forecasts and the past consensus estimates. In turn, this is an advantage of the model setup.

E Latency of Information Structures

E.1 Motivation

In the context of CG-type (2015) analyses, the standard theoretical models of information rigidity rely on information structure where the forecasted variable is a latent state. Besides it being intuitively unnatural, this approach to a forecasting problem can appear odd, given that all the reduced-form evidence used to highlight regularities in the production of forecasts has the *observed* variable on the left hand side of the central regressions:

$$x_{t+h} - \hat{x}^i_{t+h|t} = \alpha + \beta \cdot \left(\hat{x}^i_{t+h|t} - \hat{x}^i_{t+h|t-1}\right) + u_t$$

Conceptually, it does not seem reasonable to study agents' information and rationality with a regression on *realized forecast errors* if we do not let them observe (like the econometrician) the realization of the variables they forecast. Indeed, a common heuristic used in this space to convey intuition, is an arbitrage-like argument that goes as follows: *a negative (positive) coefficient β is evidence of systematic overreaction (under-response), which is not consistent with standard paradigms like full information rational expectations. If forecasters were rational*[47]*, they would use the residual predictability embedded in β to improve on their forecasting performance until the forecast revision resulted orthogonal ($\beta = 0$) to the forecast error.* Clearly, the common implicit assumption behind these ideas is the **observability** of overreaction/under-response, which in turn implies the observability of the realization of the forecasted variable.

E.2 Latency and Recursivity

The stark implication of using a non-latent structure (i.e. forecasters observing the realization of the forecasted variable at t) is a break from the recursivity of optimal forecasting rules. The easiest way to see it is to notice the general structure of what is sometimes referred to as **observation equation**, or update:

$$\mathbb{E}(x_t|\Omega^i_t) = f(\mathbb{E}(x_t|\Omega^i_{t-1}); \Omega^i_t) \tag{42}$$

(42) describes a *recursive* relationship that relates the current expectation of x_t to a function of the *previous* expectation (*stock*) and the contents (*flow*) of information Ω^i_t. The first argument of function f can be interpreted as the *optimal* output of all past information sets $\Omega^i_{\tau<t}$ (or *prior*) and the second argument as the flow of new information revealed at time t, so that their partition represents *all* the information available in the history up to t. One can think of the function f as the solution to a general class of optimization problems, among which there is the family of stochastic filtering problems, including the famous discrete-time Gaussian case, the Kalman filter.

Now, a property of the recursivity of this class of solutions is the fact that a one-period solution (or more generally the "first" iteration of a solution) is sufficient to infer all higher-order (multi-period) expectations, e.g. solutions to the optimization problem, *given* the loss function specified. For instance, in the Kalman filter, all $h > 0$ expectations are expressed in terms of today's expectation of the current realization and then projected forward using the knowledge

[47] and exclusively driven by accuracy motives, i.e. min MSE

of the autoregressive process, as in the following, which is sometimes referred to as the **state equation**:

$$\mathbb{E}(x_{t+h}|\Omega_t^i) = \rho^h(\mathbb{E}(x_t|\Omega_t^i)) \qquad h = \{0, 1, \dots\} \tag{43}$$

Notice that (42) and (43) are very different conceptual exercises, and the output of (42) is what features on the right hand side of (43), which is in fact a logically subsequent step. However, and this is the key point of this note, notice how *observing* this period's realization (x_t) "breaks" the recursive structure, as $\mathbb{E}(x_t|\Omega_t^i)$ trivially reduces to x_t if $x_t \in \Omega_t^i$. Crucially, then, in (42) the first argument of the function completely decays, as it is "*dominated*" by the direct observation of x_t. By violating the recursive dependence in (42), I argue that recursive structures **rely** on latent processes – i.e. never observed states – and relaxing this (strong) assumption implies their redundancy:

$$\mathbb{E}(x_{t+h}|\Omega_t^i) = \rho^h x_t \tag{44}$$

E.3 Necessary Caveat

All of the above is intended in the conventional case where other signals – if any – in Ω_t^i are in the form of signals about the *present* (e.g. $y_t = x_t + \omega_t^i$). This is important because, if there exists at least a signal about the *future* (e.g. $\Omega_t^i \supset b_t^i = x_{t+1} + \psi_t^i$), then the structure of the problem is altered in both latent and non-latent cases: (i) if x_t was directly observed, then (3) would be suboptimally neglecting information in b_t^i, which contains data about the future; (ii) on the other hand, if x_t was latent, then the system (1)-(2) could not possibly be expressing the optimal expectation of x_{t+h}, as $\mathbb{E}(x_t|\Omega_t^i)$ would not be a sufficient statistic anymore to infer higher-order horizons and (2) would be consequently neglecting information about the future (b_t^i)[48].

E.4 Application 1: Contemporaneous signals

An example of the above can be seen with the following information structure:

$$x_{t+1} = \rho x_t + u_{t+1} \quad u_t \sim i.i.d\left(0, \sigma_u^2\right) \tag{45}$$
$$p_t = u_t + e_t \qquad e_t \sim i.i.d.\left(0, \sigma_e^2\right) \tag{46}$$
$$\Omega_t^i = \{x_t, p_t\} \cup \Omega_{t-1}^i$$

This is a trivial case where the signal is obviously redundant and the expectation of the future coincides with the unconditional expectation:

$$\mathbb{E}(x_{t+1}|\Omega_t^i) = \rho\,\mathbb{E}(x_t|\Omega_t^i) = \rho x_t$$

[48]The deeper reason for the "failure" of the Kalman-like structure is the fact that, while in the baseline scenario *all* dynamics (information about the future) were embodied by the autoregressive process of x_t, in this alternative setup the signals in Ω_t^i themselves convey information about the future, so that an equation like the state equation – exclusively function of the 0^{th}-order expectation – is unsuited to express an optimal forecasting rule.

E.5 Application 2: Forward signals

An example of the above can be seen with the following information structure:

$$x_{t+1} = \rho x_t + u_{t+1} \quad u_t \sim i.i.d.\left(0, \sigma_u^2\right) \tag{47}$$
$$g_t = u_{t+1} + e_t \quad e_t \sim i.i.d.\left(0, \sigma_e^2\right) \tag{48}$$
$$\Omega_t^i = \{x_t, p_t\} \cup \Omega_{t-1}^i$$

First, notice that if signal p_t is about tomorrow's innovation or tomorrow's realization of x is irrelevant: if $p_t = x_{t+1} + e_t$, then $p_t = \rho x_t + u_{t+1} + e_t$, hence $p_t - \rho x_t = u_{t+1} + e_t$.

It is trivial to see that (43) does not hold:

$$\mathbb{E}(x_{t+1}|\Omega_t^i) \neq \rho\,\mathbb{E}(x_t|\Omega_t^i) = \rho x_t$$

because

$$\mathbb{E}(x_{t+1}|\Omega_t^i) = \mathbb{E}(\rho x_t + u_{t+1}|\Omega_t^i) = \rho x_t + \mathbb{E}(u_{t+1}|\Omega_t^i) = \rho x_t + p_t$$

where the third equality may be seen as a degenerate case of (42) with a zero-weight on the prior.

Notice also that in both applications above, the expectations of the future are exclusively a function of *current* information, i.e. they are exclusively Bayesian posteriors.

E.6 Generalization or Discontinuity?

One could argue that a recursive structure like the Kalman filter is a generalization of a case where the observation of x_t is not perfect. That is, if we consider the latency from a *continuous* perspective, perfectly observing today's realization corresponds to a case where there exists a signal $y_t = x_t + \omega_t$ with $\sigma_\omega^2 \to 0$. Then, the following structure should hold:

$$\mathbb{E}(x_t|\Omega_t^i) = \mathbb{E}(x_t|\Omega_{t-1}^i) + K(y_t - \mathbb{E}(x_t|\Omega_{t-1}^i)) \tag{49}$$

Where $K \to 1$ as y_t is *almost* fully revealing. In fact, $K = \frac{\sigma_\omega^{-2}}{\sigma_\omega^{-2}+\sigma_u^{-2}}$, which goes to 1 as σ_ω^2 approaches zero. Then, $\mathbb{E}(x_t|\Omega_t^i) \to y_t$.

Now, say that there exists an additional signal $p_t = x_{t+1} + e_t$. Ultimately, p_t is a joint signal of both x_t (almost perfectly observed) and u_{t+1}. If we focus on the nowcast of the *almost perfectly observed* x_t, then:

$$\begin{aligned}\mathbb{E}(x_t|\Omega_t^i) &= \mathbb{E}(x_t|\Omega_{t-1}^i) + K_y(y_t - \mathbb{E}(x_t|\Omega_{t-1}^i)) + K_g(p_t - \mathbb{E}(x_t|\Omega_{t-1}^i))\\ &\approx y_t\end{aligned}$$

Once again the expectation approaches y_t, as $K_y \to 1$ and $K_g \to 0$. Essentially, as long as y_t is almost perfect ($\sigma_w \to 0$) and p_t is not ($\sigma_{u,e} > 0$), the latter's information on x_t is completely ignored.

Notice though that the second "arm" of the Kalman-like structure (as in (43)) does *not* hold, as

p_t *is* useful in forecasting x_{t+1}, but is not utilized in a state equation that only projects forward the nowcast:

$$\mathbb{E}(x_{t+1}|\Omega_t^i) \neq \rho\,\mathbb{E}(x_t|\Omega_t^i) \approx \rho y_t$$

because

$$\mathbb{E}(x_{t+1}|\Omega_t^i) = \mathbb{E}(\rho x_t + u_{t+1}|\Omega_t^i) \approx \rho y_t + p_t$$

as indeed there are two sources of information to predict x_{t+1}: the "autoregressiveness" of the process x (for which we have an almost perfect signal) *and* the exogenous signal p_t.

F Forecaster-by-Forecaster Regression

In the following I report a sample of forecaster-by-forecaster regressions, that are obviously less precise than the pooled estimates presented in the main body. The interpretation of the regressions does not change, but from this exercise we get as many estimates as the number of forecasters in the data. Nonetheless, the estimated β^i coefficients mostly preserve the negative sign, in most instances both when considering the median and the mean. I also report the box plots of the coefficients estimated on eight macroeconomic variables from the SPF to provide further insight on their distribution.

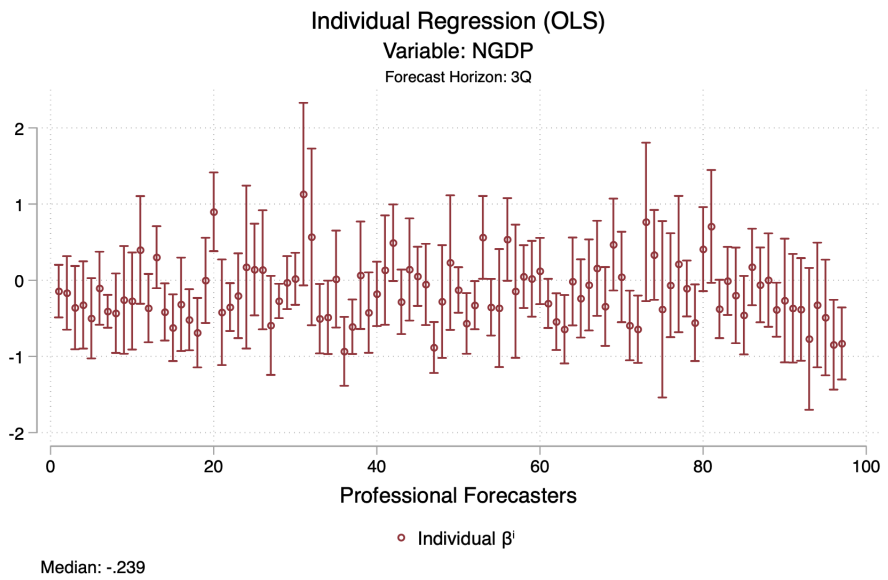

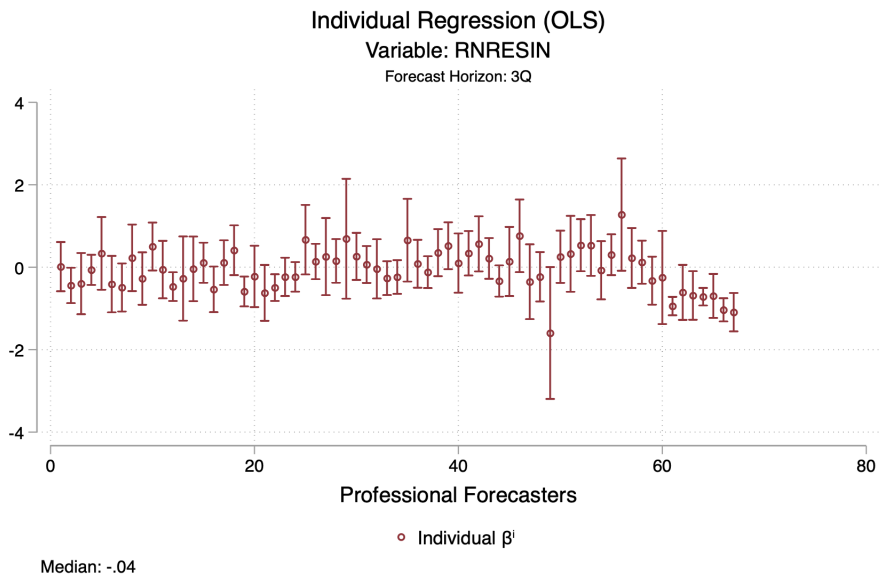

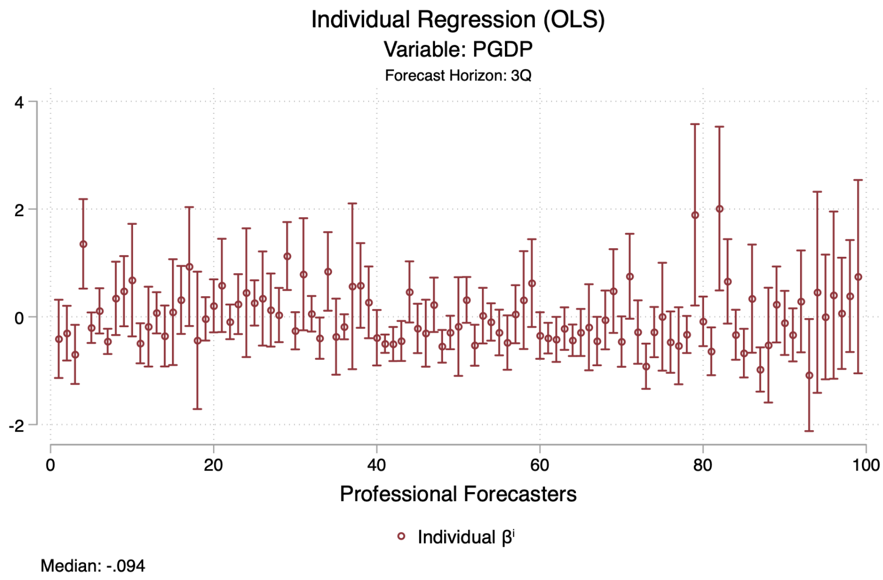
Individual Regression (OLS)
Variable: PGDP
Forecast Horizon: 3Q
4
2
0
-2
0
20
40
60
80
100
Professional Forecasters
Individual β^i
Median: -.094

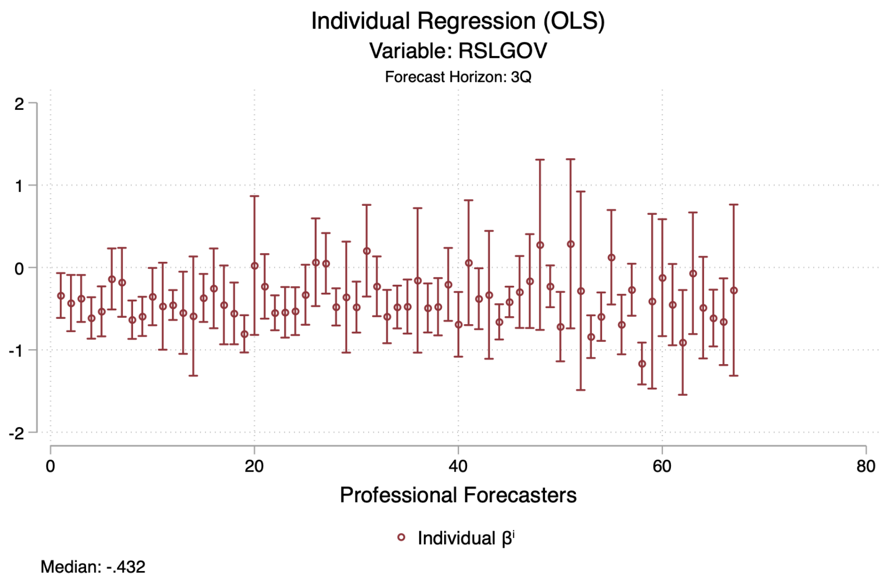
Individual Regression (OLS)
Variable: RSLGOV
Forecast Horizon: 3Q
2
1
0
-1
-2
0
20
40
60
80
Professional Forecasters
Individual β^i
Median: -.432

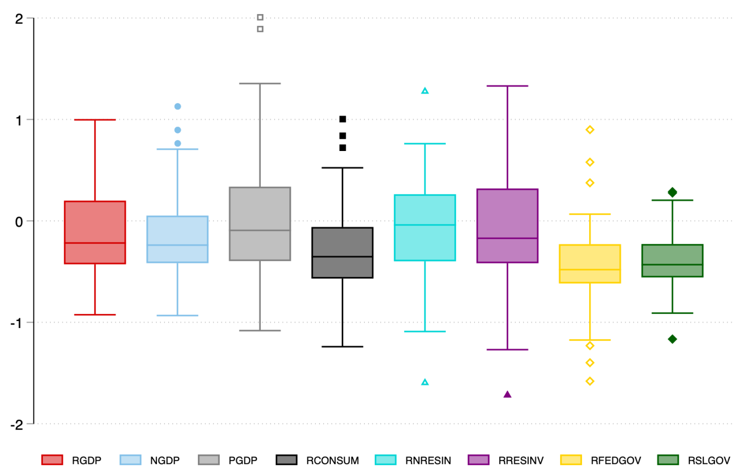
2
1
0
-1
-2
RGDP
NGDP
PGDP
RCONSUM
RNRESIN
RRESINV
RFEDGOV
RSLGOV

G Two-types & multi-period signals

This section delineates a version of the model with a richer information structure where the forecasters are of two types based on the private information they receive.
Formally, the model is as follows:

$$x_{t+1} = \rho x_t + u_{t+1}$$

$$u_{t+1} = p_t + \sum_i^2 y_t^i + \sum_i^2 r_{t-1}^i + \epsilon_{t+1} \tag{50}$$

$$p_t \sim \mathcal{N}\left(0, \sigma_p^2\right) \quad y_t^i \sim \mathcal{N}\left(0, \sigma_s^2\right) \quad r_t \sim \mathcal{N}\left(0, \sigma_r^2\right) \quad \epsilon_{t+1} \sim \mathcal{N}\left(0, \sigma_\epsilon^2\right) \tag{51}$$

The innovation depicted in (**??**) is then a function of the usual p_t (observed by everyone) and ϵ_{t+1} (observed by no one), plus two-types of signals : y_t^i is information about a one-period ahead realization, while r_t^i represents a two-periods signal. We can then assume that forecasters observe the consensus forecasts $\bar{\hat{x}}_{t+1|t}$ and $\bar{\hat{x}}_{t+2|t}$ with a one-period lag, that is, in $t+1$[49]. Then, the information set of forecaster i at time t is given by:

$$\Omega_t^i = \{x_t, p_t, y_t^i, r_{t-1}^i, r_{t-1}^{-i}, \bar{\hat{x}}_{t|t-1}, \bar{\hat{x}}_{t+1|t-1}\}$$

The optimization problem is unchanged:

$$\text{FOC}_{h=1} : \hat{x}_{t+1|t}^i = \frac{1}{1+\phi}\mu_{t+1|t}^i + \frac{\phi}{1+\phi}\mathbb{E}_i\left(\bar{\hat{x}}_{t+1|t}\right)$$

Where

$$\mu_{t+1|t}^i = \rho x_t + p_t + y_t^i + r_{t-1}^i + r_{t-1}^{-i} \tag{52}$$

We formulate a conjecture based on information that is public or private:

$$\hat{x}_{t+1|t}^i = \psi_0\left(\rho x_t + p_t + r_{t-1}^i + r_{t-1}^{-i}\right) + \psi_1^i\left(y_t^i\right) \tag{53}$$

Then

$$\mathbb{E}_i\left(\bar{\hat{x}}_{t+1|t}\right) = \psi_0\left(\rho x_t + p_t + r_{t-1}^i + r_{t-1}^{-i}\right) + \psi_1^i p^i\left(y_t^i\right)$$

Where p^i represents the relative share of i-type forecasters. Then the F.O.C. becomes

$$\hat{x}_{t+1|t}^i = \frac{1}{1+\phi}(\rho x_t + p_t + y_t^i + r_{t-1}^i + r_{t-1}^{-i}) + \frac{\phi}{1+\phi}\left(\psi_0\left(\rho x_t + p_t + r_{t-1}^i + r_{t-1}^{-i}\right) + \psi_1^i p^i y_t^i\right)$$

Matching coefficients delivers

$$\psi_0 = 1 \quad \psi_1^i = \frac{1}{1+\phi\left(1-p^i\right)}$$

[49] It can be proven that this assumption is redundant for the purpose of the exercise. Observing past consensuses allows (provided that the number of types does not exceed 2) for the extraction of past private signals, which can be used in updating own forecasts. However, this will result only in minor quantitative changes in the computation for the β^i and β^c.

Analogously, the two-periods forecasts is:

$$\text{FOC}_{h=2} : \hat{x}^i_{t+2|t} = \frac{1}{1+\phi}\mu^i_{t+2|t} + \frac{\phi}{1+\phi}\mathbb{E}_i\left(\bar{\hat{x}}_{t+2|t}\right)$$

Where

$$\mu^i_{t+2|t} = \rho^2 x_t + \rho(p_t + y^i_t + r^i_{t-1} + r^{-i}_{t-1}) + r^i_t \tag{54}$$

Once again, our conjecture is:

$$\hat{x}^i_{t+1|t} = \theta_0\left(\rho^2 x_t + \rho(p_t + r^i_{t-1} + r^{-i}_{t-1})\right) + \theta^i_1\left(\rho y^i_t + r^i_t\right) \tag{55}$$

and then

$$\mathbb{E}_i\left(\bar{\hat{x}}_{t+2|t}\right) = \theta_0\left(\rho^2 x_t + \rho(p_t + r^i_{t-1} + r^{-i}_{t-1})\right) + \theta^i_1 p^i\left(\rho y^i_t + r^i_t\right)$$

Hence

$$\begin{aligned} \hat{x}^i_{t+2|t} = & \frac{1}{1+\phi}\left(\rho^2 x_t + \rho(p_t + y^i_t + r^i_{t-1} + r^{-i}_{t-1}) + r^i_t\right) + \\ & + \frac{\phi}{1+\phi}\left(\theta_0\left(\rho^2 x_t + \rho(p_t + r^i_{t-1} + r^{-i}_{t-1})\right) + \theta^i_1 p^i\left(\rho y^i_t + r^i_t\right)\right) \end{aligned} \tag{56}$$

Matching coefficients delivers

$$\theta_0 = 1 \quad \theta^i_1 = \frac{1}{1+\phi\left(1-p^i\right)}$$

Finally, we have two expression for a 1 and 2-periods ahead announced forecasts:

$$\hat{x}^i_{t+1|t} = \rho x_t + p_t + r^i_{t-1} + r^{-i}_{t-1} + \frac{1}{1+\phi\left(1-p^i\right)}y^i_t \tag{57}$$

$$\hat{x}^i_{t+2|t} = \rho^2 x_t + \rho\left(p_t + r^i_{t-1} + r^{-1}_{t-1}\right) + \frac{1}{1+\phi\left(1-p^i\right)}\left[\rho y^i_t + r^i_t\right] \tag{58}$$

What does this richer framework imply for our individual regression?
Endowed with the two expressions above, I construct the model implied FE^i_{t+1} and FR^i_t:

$$\begin{aligned} FE^i_{t+1} &= x_{t+1} - \hat{x}^i_{t+1|t} = \rho x_t + u_{t+1} - \rho x_t - p_t - r^i_{t-1} - r^{-i}_{t-1} - \frac{1}{1+\phi\left(1-p^i\right)}y^i_t \\ &= p_t + \sum_i^2 y^i_t + \sum_i^2 r^i_{t-1} + \varepsilon_{t+1} - p_t - r^i_{t-1} - r^{-i}_{t-1} - \frac{1}{1+\phi\left(1-p^i\right)}y^i_t \\ &= \varepsilon_{t+1} + y^{-i}_t + \frac{\phi\left(1-p^i\right)}{1+\phi\left(1-p^i\right)}y^i_t \end{aligned} \tag{59}$$

$$
\begin{aligned}
FR_t^i = \hat{x}^i_{t+1|t} - \hat{x}^i_{t+1|t-1} &= \rho x_t + p_t + r^i_{t-1} + r^{-i}_{t-1} + \frac{1}{1+\phi\left(1-p^i\right)} y^i_t - \\
&- \rho^2 x_{t-1} - \rho\left(p_{t-1} + r^i_{t-2} + r^{-i}_{t-2}\right) - \frac{1}{1+\phi\left(1-p^i\right)}\left(\rho y^i_{t-1} + r^i_{t-1}\right) \\
&= \rho\left(p_{t-1} + \sum_i^2 y^i_{t-1} + \sum_i^2 r^i_{t-2} + \epsilon_t\right) + p_t + \sum_i^2 r^i_{t-1} + \frac{1}{1+\phi\left(1-p^i\right)} y^i_t - \\
&- \rho\left(p_{t-1} + \sum_i^2 r^i_{t-2}\right) - \frac{1}{1+\phi\left(1-p_i\right)}\left(\rho y^i_{t-1} + r^i_{t-1}\right) \\
&= \rho y^{-i}_{t-1} + \frac{\rho\phi\left(1-p^i\right)}{1+\phi\left(1-p^i\right)} y^i_{t-1} + \rho\epsilon_t + p_t + \frac{\phi\left(1-p^i\right)}{1+\phi\left(1-p^i\right)} r^i_{t-1} + r^{-i}_{t-1} + \frac{1}{1+\phi\left(1-p^i\right)} y^i_t
\end{aligned}
\tag{60}
$$

Hence

$$
\begin{aligned}
\beta^i &= \frac{\text{Cov}\left(\epsilon_{t+1} + y^{-i}_t + \frac{\phi\left(1-p^i\right)}{1+\phi\left(1-p^i\right)} y^i_t, \rho y^{-i}_{t-1} + \frac{\rho\phi\left(1-p^i\right)}{1+\phi\left(1-p^i\right)} y^i_{t-1} + \rho\epsilon_t + p_t + \frac{\phi\left(1-p^i\right)}{1+\phi\left(1-p^i\right)} r^i_{t-1} + r^{-i}_{t-1} + \frac{1}{1+\phi\left(1-p^i\right)} y^i_t\right)}{\text{Var}\left(\rho y^{-i}_{t-1} + \frac{\rho\phi\left(1-p^i\right)}{1+\phi\left(1-p^i\right)} y^i_{t-1} + \rho\epsilon_t + p_t + \frac{\phi\left(1-p^i\right)}{1+\phi\left(1-p^i\right)} r^i_{t-1} + r^{-i}_{t-1} + \frac{1}{1+\phi\left(1-p^i\right)} y^i_t\right)} \\
&= \frac{QA\sigma_y^2}{\sigma_y^2\left[\rho^2\left(Q^2+1\right) + A^2\right] + \sigma_r^2\left(Q^2+1\right) + \sigma_p^2 + \rho\sigma_\epsilon^2}
\end{aligned}
\tag{61}
$$

Comparing (61) with the body, we observe that they are nearly equivalent, with the difference of a richer forecast revision structure – and with it, an additional parameter (σ_r^2) to be estimated. In summary, a model featuring multi-period signals and multi-period forecasting will result essentially in the same implications of a single-period signal and a single-period forecast projected forward by multiplying it by ρ^h.

H Extended Expressions for Regression Coefficients

OLS formulas for the bivariate linear regression provide the direct mapping from the model to the empirical estimates:

$$\beta_A^p = \frac{\text{Cov}(FR_t^\tau, FE_{t+1}^\tau)\,\text{Var}(DFC_t^\tau) - \text{Cov}(DFC_t^\tau, FR_t^\tau)\,\text{Cov}\big(DFC_t^\tau, FE_{t+1}^\tau\big)}{\text{Var}(FR_t^\tau)\,\text{Var}(DFC_t^\tau) - \text{Cov}(FR_t^\tau, DFC_t^\tau)^2} \tag{62}$$

$$\gamma_A^p = \frac{\text{Cov}(DFC_t^\tau, FE_{t+1}^\tau)\,\text{Var}(FR_t^\tau) - \text{Cov}(DFC_t^\tau, FR_t^\tau)\,\text{Cov}\big(FR_t^\tau, FE_{t+1}^\tau\big)}{\text{Var}(FR_t^\tau)\,\text{Var}(DFC_t^\tau) - \text{Cov}(FR_t^\tau, DFC_t^\tau)^2} \tag{63}$$

where the new terms I need an expression for are computed below:

$$\begin{aligned}
\text{Cov}\left(DFC_t^i, FE_{t+1}^i\right) &= \text{Cov}\left(\frac{N-1}{N+\phi\,(N-1)}y_t^i - \frac{1}{N+\phi\,(N-1)}\sum_{j\neq i}^N y_t^j, \sum_{j\neq i}^N y_t^j + \varepsilon_{t+1} + \underbrace{\frac{\phi(N-1)}{N+\phi\,(N-1)}}_{Q} y_t^i\right) \\
&= \text{Cov}\left(\frac{Q}{\phi}y_t^i, Qy_t^i\right) - \text{Cov}\left(\frac{A}{N}\sum_{j\neq i}^N y_t^j, \sum_{j\neq i}^N y_t^j\right) \\
&= \frac{Q^2}{\phi}\sigma_y^2 - \frac{A}{N}(N-1)\sigma_y^2 = \frac{NQ^2 - A\phi(N-1)}{N\phi}\sigma_y^2 \\
\text{Cov}(DFC_t^i, FR_t^i) &= \text{Cov}\left(\frac{Q}{\phi}y_t^i - \frac{A}{N}\sum_{j\neq i}^N y_t^j, \rho\left[\varepsilon_t + Qy_{t-1}^i + \sum_{j\neq i}^N y_{t-1}^j\right] + p_t + Ay_t^i\right) \\
&= \frac{QA}{\phi}\sigma_y^2 \\
\text{Var}(DFC_t^i) &= \text{Var}\left(\frac{Q}{\phi}y_t^i - \frac{A}{N}\sum_{j\neq i}^N y_t^j\right) = \left(\frac{Q}{\phi}\right)^2\sigma_y^2 + \left(\frac{A}{N}\right)^2(N-1)\sigma_y^2 \\
&= \left(\left(\frac{Q}{\phi}\right)^2 + \left(\frac{A}{N}\right)^2(N-1)\right)\sigma_y^2
\end{aligned}$$

Hence

$$\beta_A^i = \frac{QA\sigma_y^2\left[\left(\frac{Q}{\phi}\right)^2 + \left(\frac{A}{N}\right)^2(N-1) - \sigma_y^2\left(\left(\frac{Q}{\phi}\right)^2 - \frac{A}{\phi N}(N-1)\right)\right]}{\left[\sigma_y^2\left(\rho^2Q^2 + \rho^2(N-1) + A^2\right) + \rho^2\sigma_\varepsilon^2 + \sigma_p^2\right]\left[\left(\frac{Q}{\phi}\right)^2 + \left(\frac{A}{N}\right)^2(N-1)\right] - \left(\frac{QA}{\phi}\sigma_y^2\right)^2} \tag{64}$$

$$\gamma_A^i = \frac{\left(\frac{Q^2}{\phi}\sigma_y^2 - \frac{A}{N}(N-1)\sigma_y^2\right)\left[\sigma_y^2\left(\rho^2Q^2 + \rho^2(N-1) + A^2\right) + \rho^2\sigma_\varepsilon^2 + \sigma_p^2\right] - \frac{(QA\sigma_y^2)^2}{\phi}}{\left[\sigma_y^2\left(\rho^2Q^2 + \rho^2(N-1) + A^2\right) + \rho^2\sigma_\varepsilon^2 + \sigma_p^2\right]\left[\left(\frac{Q}{\phi}\right)^2 + \left(\frac{A}{N}\right)^2(N-1)\right] - \left(\frac{QA}{\phi}\sigma_y^2\right)^2} \tag{65}$$

All four coefficients, then:

$$\beta_1^p = \frac{\frac{\phi\left(1-N^{-1}\right)}{(1+\phi(1-N^{-1}))^2}\sigma_y^2}{\sigma_y^2\left[\rho^2\left(\frac{\phi(N-1)}{N+\phi(N-1)}\right)^2+\rho^2(N-1)+\left(\frac{N}{N+\phi(N-1)}\right)^2\right]+\rho^2\sigma_\varepsilon^2+\sigma_p^2}$$

$$\beta^c = \frac{\frac{N(N-1)(1+\phi)}{(N+\phi(N-1))^2}\sigma_y^2}{\rho^2\left[\left(\frac{(N-1)(1+\phi)}{N+\phi(N-1)}\right)^2 N\sigma_y^2+\sigma_\varepsilon^2\right]+\sigma_p^2+\left(\frac{1}{N+\phi(N-1)}\right)^2 N\sigma_y^2}$$

$$\beta_{Aug}^i = \frac{\left(\frac{\phi(N-1)}{(N+\phi(N-1))^2}\right)\sigma_y^2\left[\left(\frac{(N-1)}{N+\phi(N-1)}\right)^2+\left(\frac{1}{N(N+\phi(N-1))}\right)^2(N-1)-\sigma_y^2\left(\left(\frac{(N-1)}{N+\phi(N-1)}\right)^2-\frac{N-1}{(N+\phi(N-1))\phi N}\right)\right]}{\left[\sigma_y^2\left(\rho^2\left(\frac{\phi(N-1)}{N+\phi(N-1)}\right)^2+\rho^2(N-1)+\left(\frac{1}{N+\phi(N-1)}\right)^2\right)+\rho^2\sigma_\varepsilon^2+\sigma_p^2\right]\left[\left(\frac{(N-1)}{N+\phi(N-1)}\right)^2+\left(\frac{1}{N(N+\phi(N-1))}\right)^2(N-1)\right]-\left(\frac{N-1}{(N+\phi(N-1))^2}\sigma_y^2\right)^2}$$

$$\gamma_{Aug}^i = \frac{\left(\phi\left(\frac{N-1}{N+\phi(N-1)}\right)^2-\frac{N-1}{N(N+\phi(N-1))}\right)\sigma_y^2\left[\sigma_y^2\left(\rho^2\left(\frac{\phi(N-1)}{N+\phi(N-1)}\right)^2+\rho^2(N-1)+\left(\frac{1}{N+\phi(N-1)}\right)^2\right)+\rho^2\sigma_\varepsilon^2+\sigma_p^2\right]-\left(\frac{N-1}{(N+\phi(N-1))^2}\sigma_y^2\right)^2}{\left[\left(\left(\frac{\rho\phi(N-1)}{N+\phi(N-1)}\right)^2+\rho^2(N-1)+\left(\frac{1}{N+\phi(N-1)}\right)^2\right)\sigma_y^2+\rho^2\sigma_\varepsilon^2+\sigma_p^2\right]\left[\left(\frac{(N-1)}{N+\phi(N-1)}\right)^2+\left(\frac{1}{N(N+\phi(N-1))}\right)^2(N-1)\right]-\left(\frac{N-1}{(N+\phi(N-1))^2}\sigma_y^2\right)^2}$$

I Alternative Identification Strategies

At its core, the identification approach is a moment estimation similar to GMM and Minimum Distance estimation, giving equal emphasis to the selected moment conditions. These moment conditions can encompass any data feature, as the model comprehensively describes the data generating process. In the following sections, I will detail results for other specifications, altering one or more moment conditions. The rationale for choosing specific moments centered on their informativeness with respect to the target parameters and their relevance to the forecasting situations discussed.

I.1 Alternative #1: *bbVg*

bbVg stands for: β^p, β^c, $\text{Var}(u_t^p)$, γ_A^p. When these are the targeted moments, results are plotted below both in the unrestricted ($\phi \neq 0$) and restricted case ($\phi = 0$):

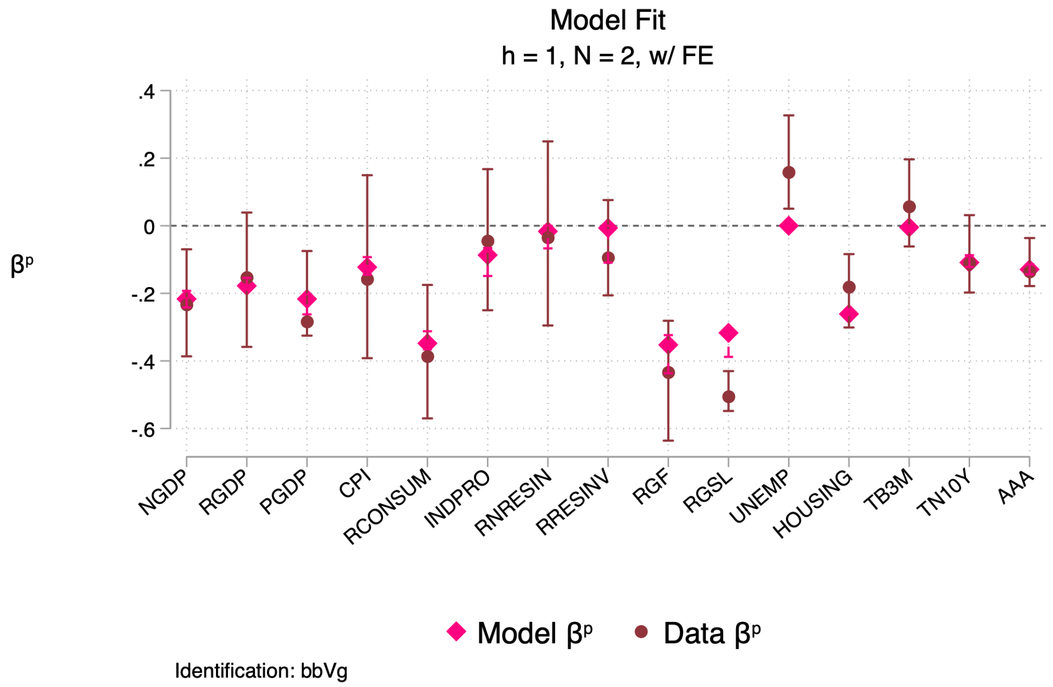

Figure 16: Model fit: Individual coefficients

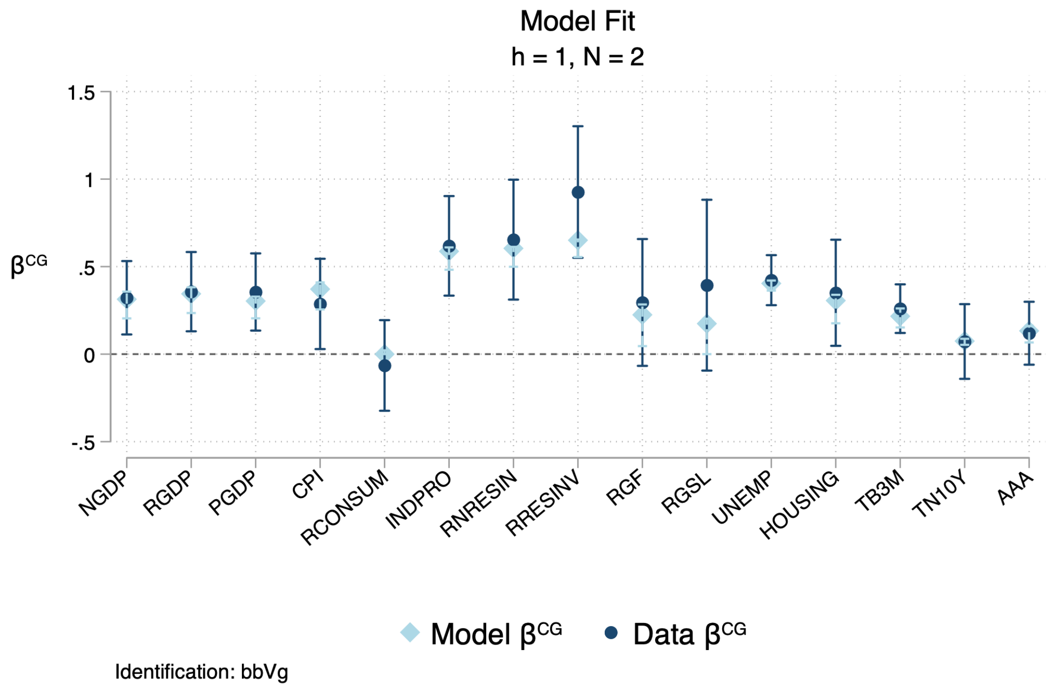

Figure 17: Model fit: Consensus coefficients

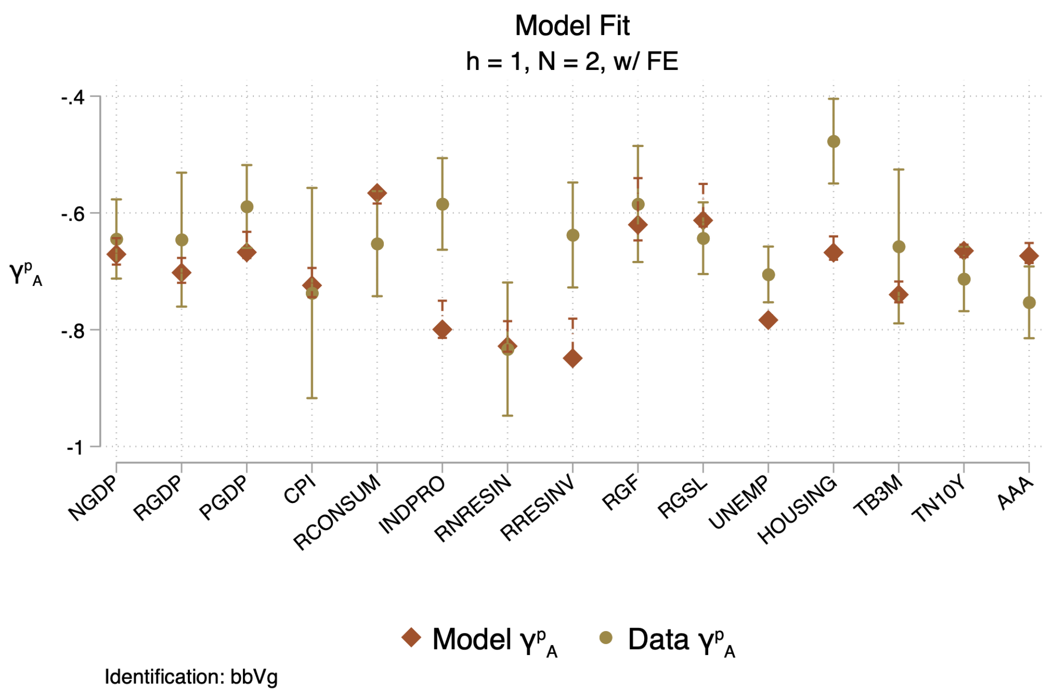

Figure 18: Model fit: Augmented γ coefficients

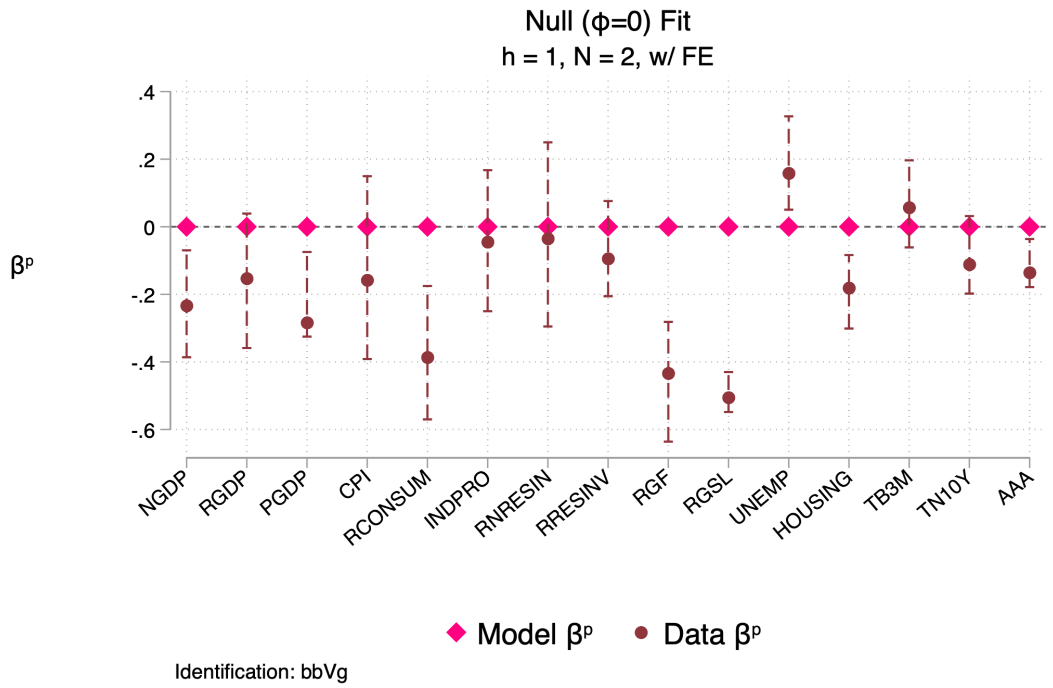

Figure 19: Null fit: Individual coefficients

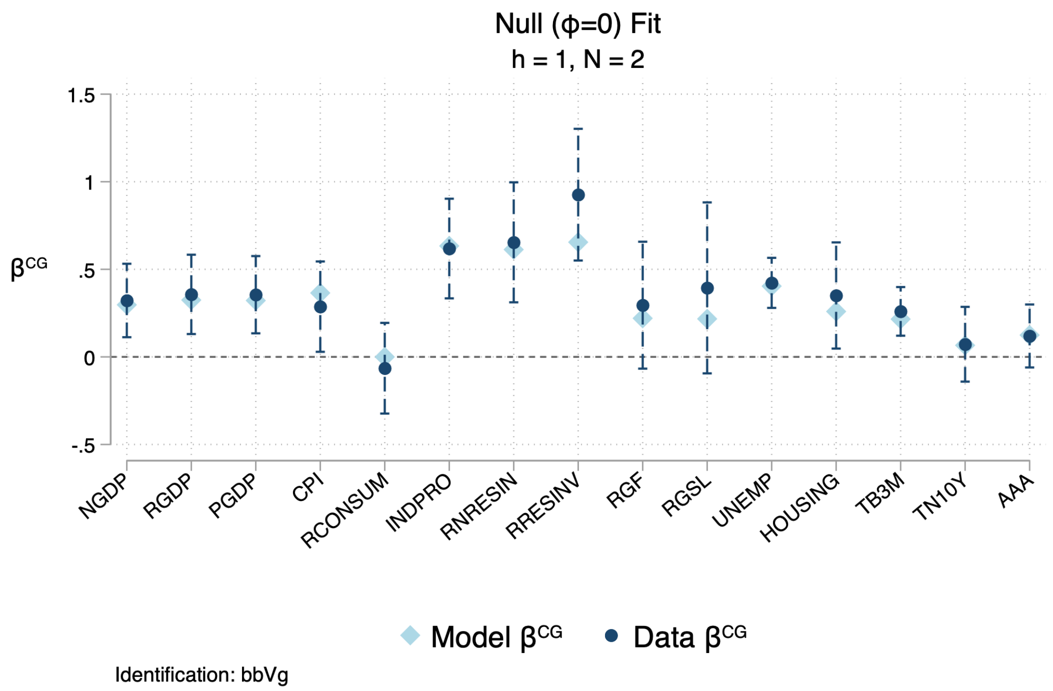

Figure 20: Null fit: Consensus coefficients

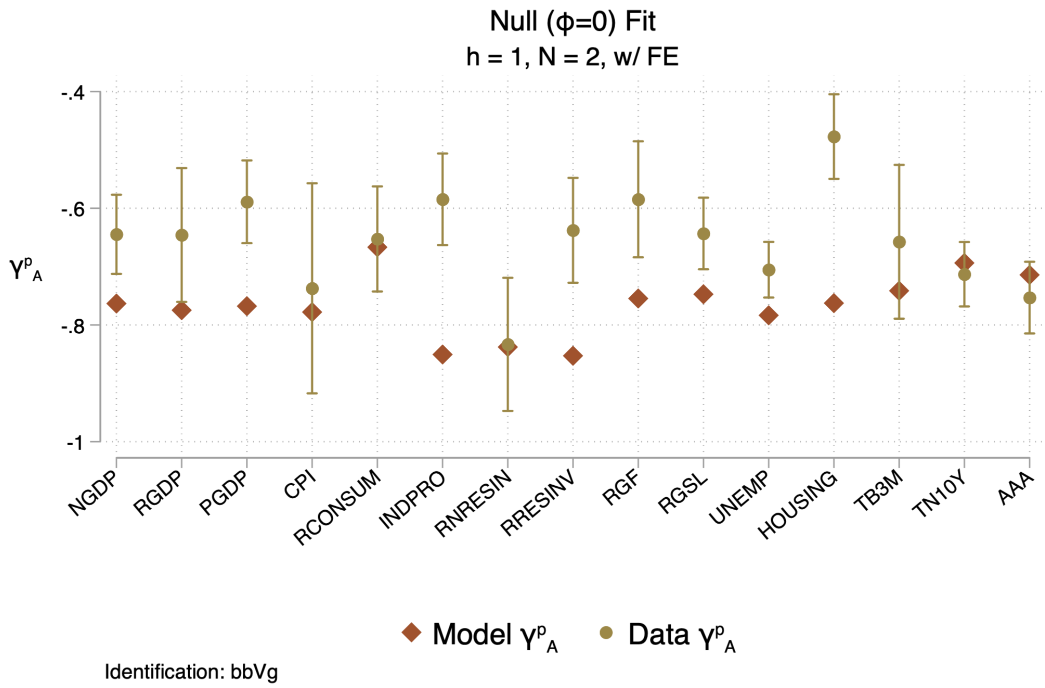

Figure 21: Null fit: Augmented γ coefficients

I.2 Alternative #2: *bbVV*

bbVV stands for: β^p, β^c, Var(u^p), Var(u^{CG}). When these are the targeted moments, results are plotted below both in the unrestricted ($\phi \neq 0$) and restricted case ($\phi = 0$):

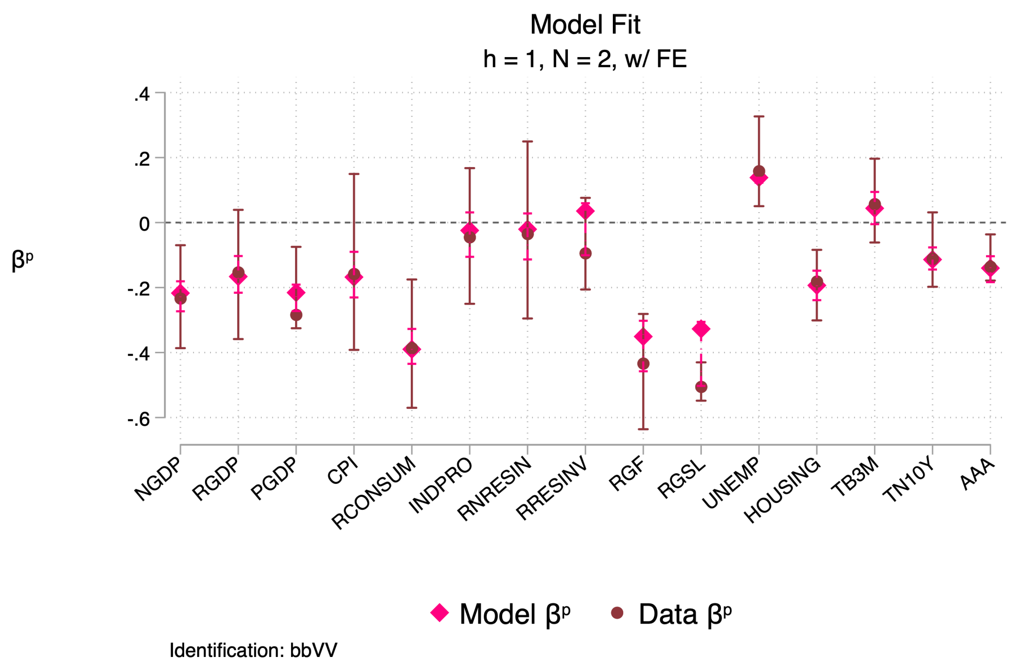

Figure 22: Model fit: Individual coefficients

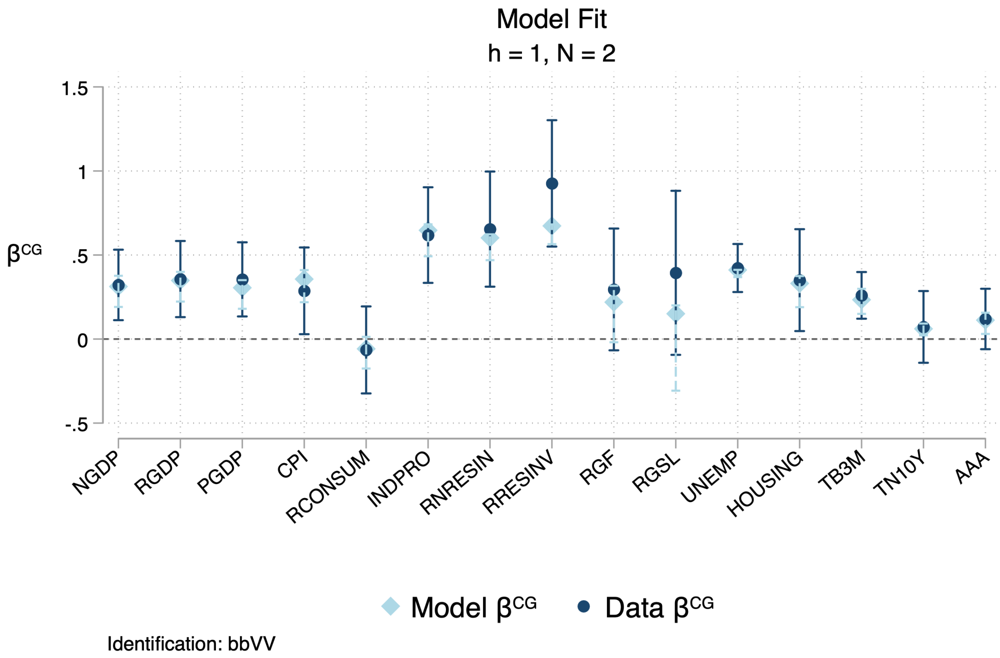

Figure 23: Model fit: Consensus coefficients

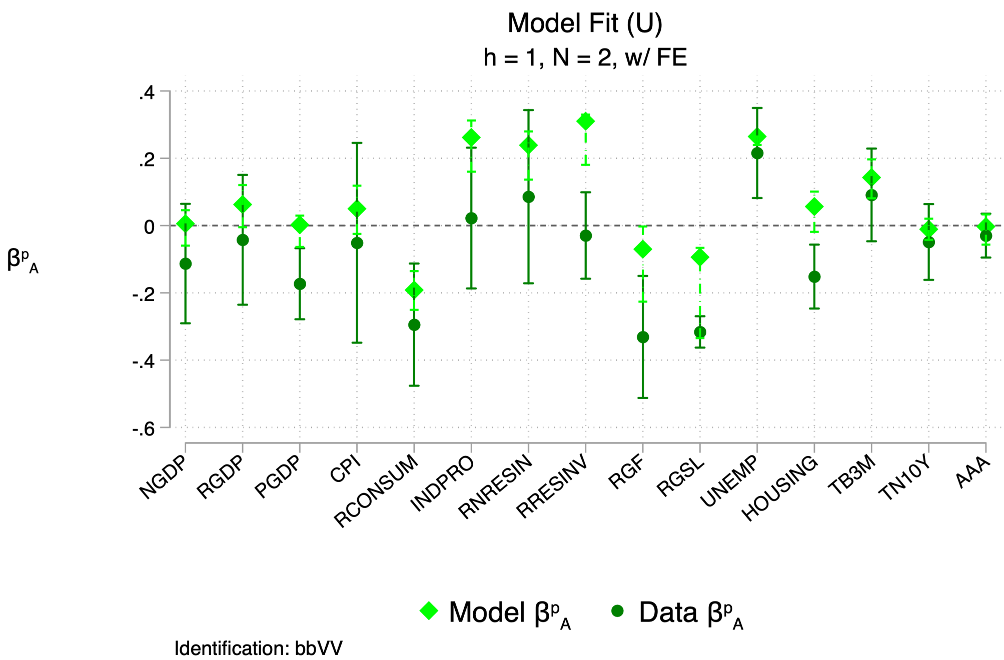

Figure 24: Model fit: Augmented β coefficients

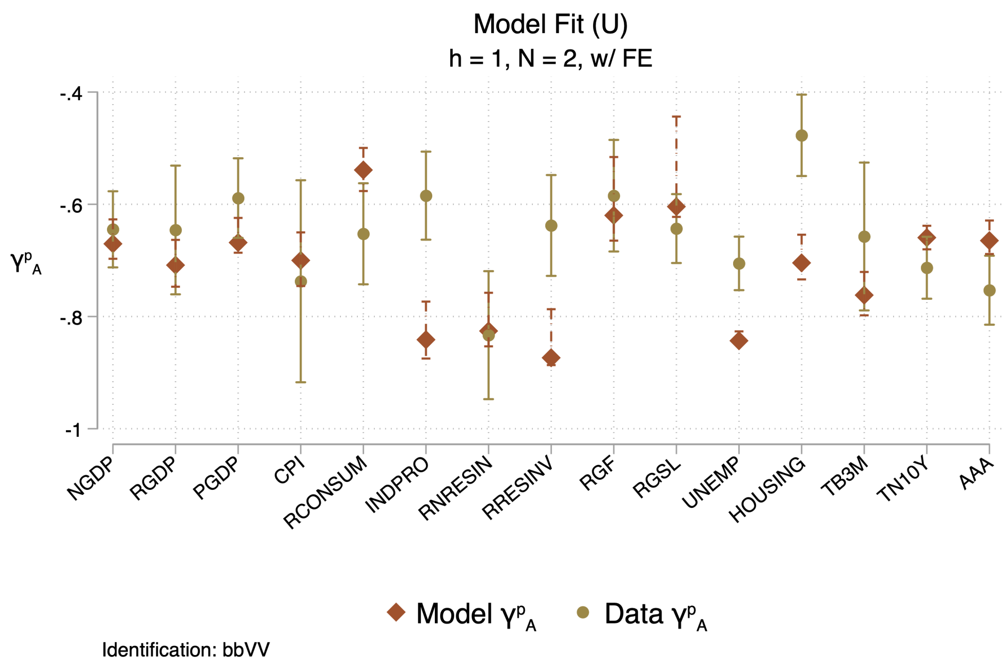

Figure 25: Model fit: Augmented γ coefficients

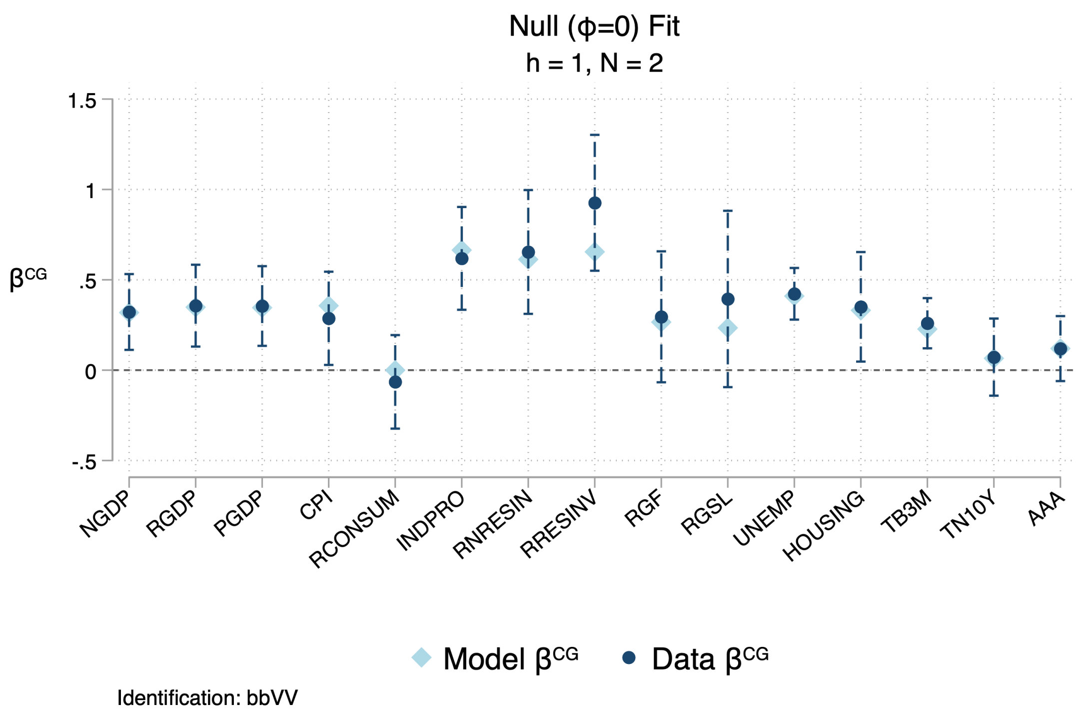

Figure 26: Model fit: Consensus coefficients

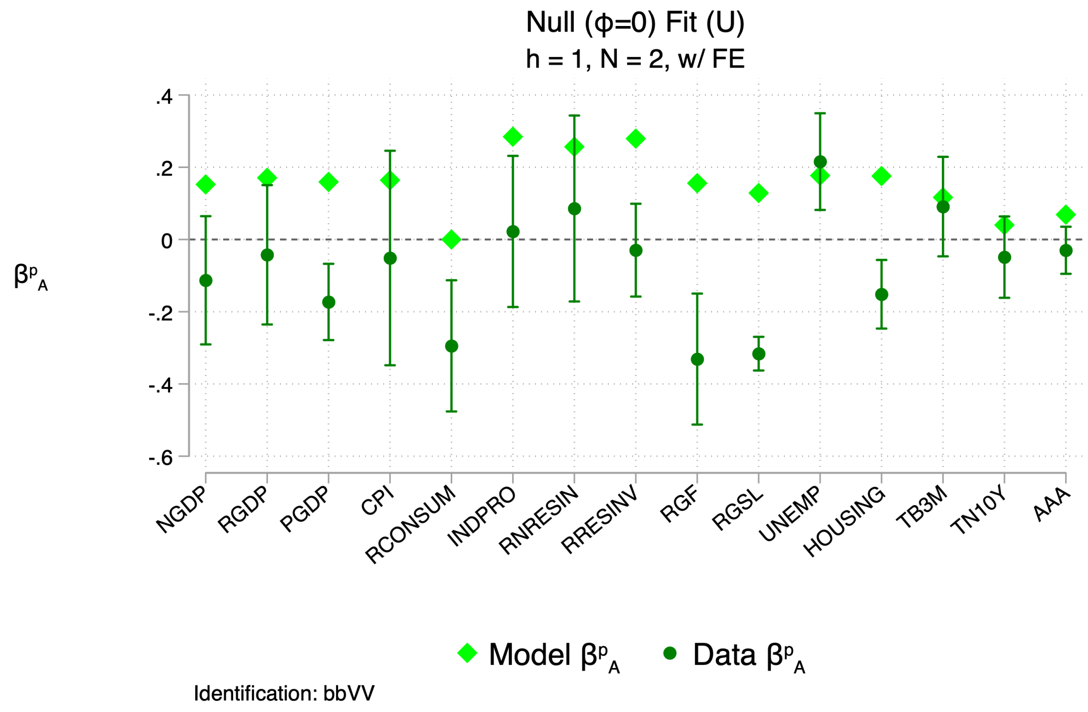

Figure 27: Model fit: Augmented β coefficients

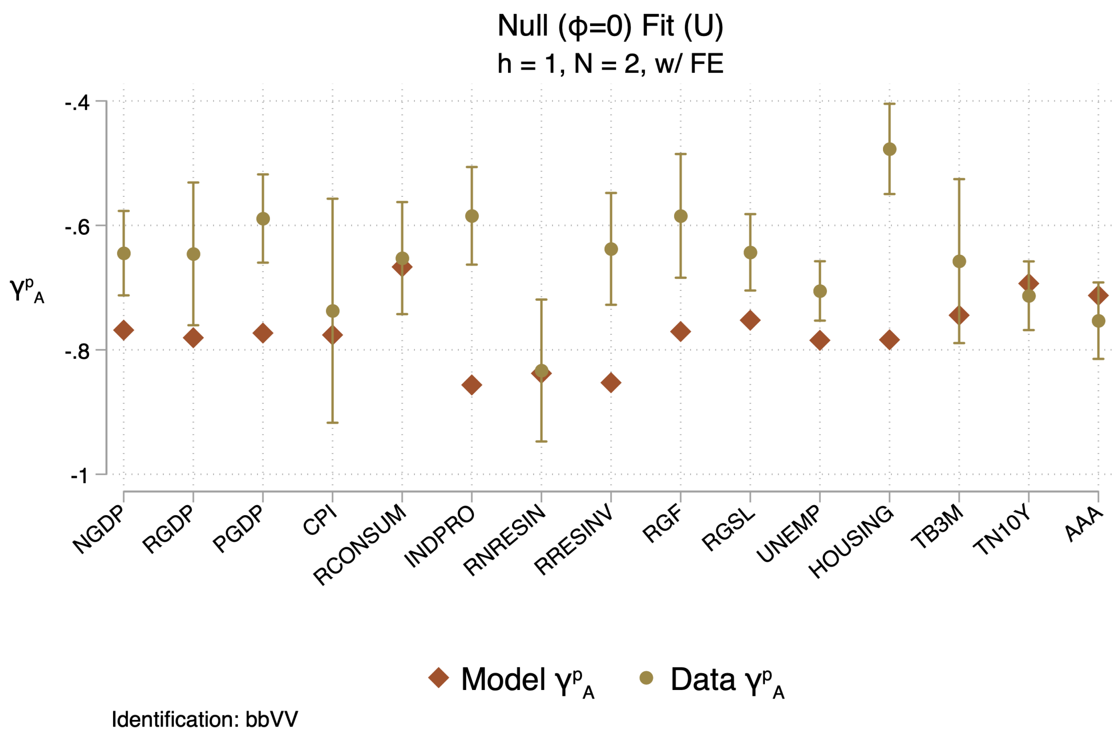

Figure 28: Model fit: Augmented γ coefficients

Table 5: Identified model parameters under *bbVV* strategy

Variable	ϕ	σ_y	σ_ε	σ_p
NGDP	-0.653	0.147	0.001	0.001
RGDP	-0.515	0.133	0.051	0.004
PGDP	-0.661	0.081	0.000	0.000
CPI	-0.539	0.098	0.021	0.005
RCONSUM	-1.060	0.086	0.017	0.003
INDPROD	-0.071	0.353	0.001	0.000
RNRESIN	-0.066	0.455	0.001	0.000
RRESINV	0.113	0.713	0.000	0.000
RGF	-0.775	0.360	0.088	0.001
RGSL	-0.847	0.107	0.000	0.000
UNEMP	0.924	0.562	0.002	0.002
HOUSING	-0.545	0.138	0.076	0.003
tb3m	0.337	0.565	0.420	0.004
tn10y	-0.730	0.323	0.466	0.005
AAA	-0.679	0.462	0.463	0.004

Note: Identification by targeting β^p, β^c, β^p_A, γ^p_A. Results depicted are for $N = 2$, $h = 1$.

I.3 Alternative #3: *bbbgVV*: $n = 6, k = 5$: Estimating N

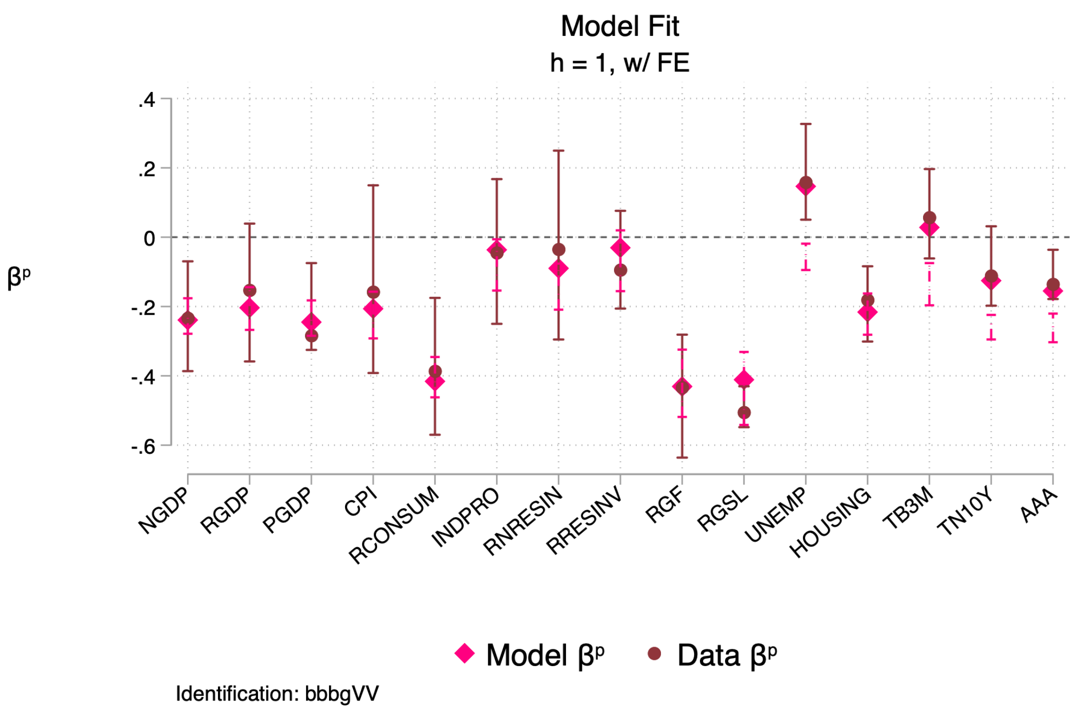

Figure 29: Model fit: Individual β coefficients

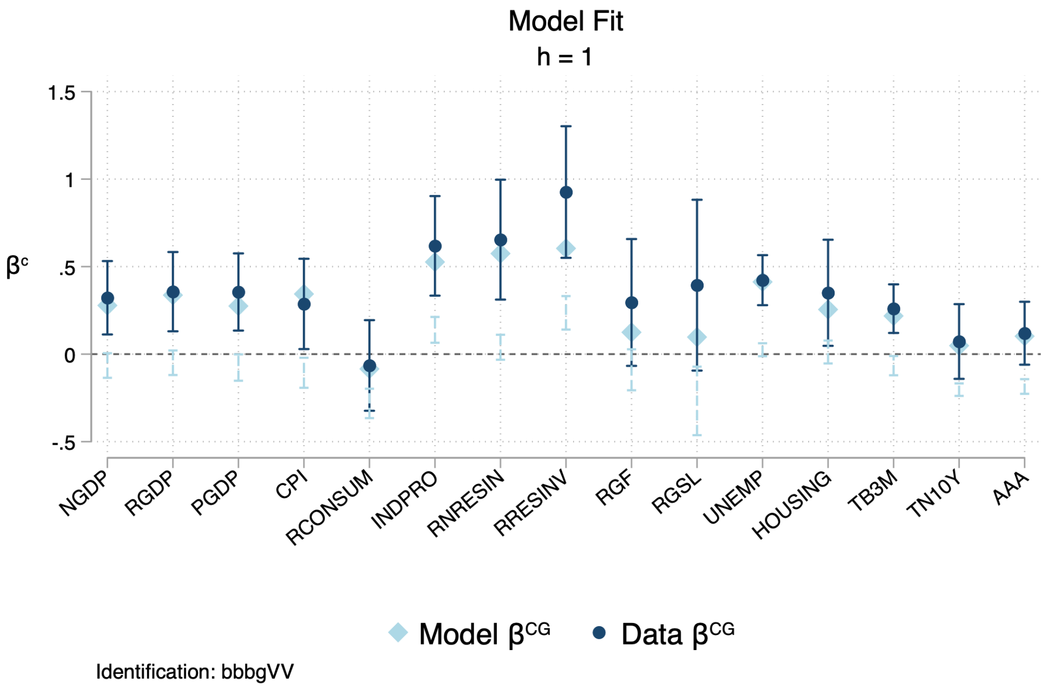

Figure 30: Model fit: Consensus β coefficients

Table 6: Identified model parameters – ID strategy: *bbbgVV*

Variable	N	ϕ	$\frac{\sigma_y}{\sum\sigma}$	$\frac{\sigma_p}{\sum\sigma}$	$\frac{\sigma_\varepsilon}{\sum\sigma}$
NGDP	2.0258	-0.6579	0.9959	0.0021	0.0021
RGDP	1.9306	-0.5250	0.8697	0.0163	0.1140
PGDP	2.0835	-0.6931	0.9975	0.0013	0.0013
CPI	2.3111	-0.6501	0.7258	0.0468	0.2274
RCONSUM	2.3926	-1.1509	0.5784	0.0450	0.3766
INDPROD	1.7164	-0.0896	0.9982	0.0008	0.0010
RNRESIN	2.1078	-0.2531	0.9977	0.0009	0.0014
RRESINV	1.8998	-0.0807	0.9991	0.0004	0.0005
RGF	2.1364	-0.8703	0.9895	0.0011	0.0094
RGSL	2.4548	-1.0163	0.9969	0.0010	0.0021
UNEMP	1.7263	0.7772	0.7827	0.0105	0.2068
HOUSING	1.7057	-0.4774	0.9962	0.0016	0.0022
tb3m	1.7563	0.1673	0.5501	0.0032	0.4467
tn10y	2.1167	-0.7673	0.3849	0.0073	0.6078
AAA	2.1923	-0.7258	0.4759	0.0046	0.5195

Note: identification achieved targeting β^p, β^c, β^p_A, γ^p_A, $\sigma^2_{u_p}$, $\sigma^2_{u_c}$. Depicted, $h = 1$.

J Information Structures: an Equivalence

The model proposed in this paper differs from the standard characterization of imperfect information. In particular, the information structure advanced is the following:

$$x_{t+1} = \rho x_t + u_{t+1}$$
$$u_{t+1} = p_t + \sum_{\tau}^{N} y_t^\tau + \varepsilon_{t+1} \tag{66}$$
$$p_t \sim i.i.d.\left(0, \sigma_p^2\right) \quad y_t^\tau \sim i.i.d.\left(0, \sigma_y^2\right) \quad \varepsilon_{t+1} \sim i.i.d.\left(0, \sigma_\varepsilon^2\right)$$
$$\Omega_t^i = \{x_t, p_t, y_t^\tau\} \cup \Omega_{t-1}^i \tag{67}$$

with p_t representing publicly observed information revealed at time t, and y_t^τ representing privately observed information revealed to type-τ forecasters exclusively. ε_{t+1}, on the contrary, represents the portion of tomorrow's innovation u_{t+1} which is purely stochastic and unforecastable. Notice that this is an additional, novel element that has no equivalent in the noisy information setting described below. This note uses the fully-fledged model to build a representation equivalence, but it is useful to keep in mind that a model without ε_{t+1} would be conceptually closer to the spirit of this mapping. To build intuition, consider the following limit case: even if $N = 1$ (i.e. there exists no information asymmetry), forecasters in this framework will necessarily make a non-null forecast error, due to the presence of an unpredictable component; on the contrary, in the noisy information world, a perfectly precise signal reveals the state, the true value of future x_{t+1}.

The traditional characterization of frictional information in economic modeling goes back to Grossman and Stiglitz (1980), but in macroeconomics is often referred to as *noisy information* and the seminal contribution is due to Woodford (2001). The structure of information is as follows[50]:

$$x_{t+1} = \rho x_t + u_{t+1} \qquad u_{t+1} \sim \mathcal{N}\left(0, \sigma_u^2\right)$$
$$\tilde{p}_t = u_{t+1} + \tilde{e}_t \qquad \tilde{e}_t \sim \mathcal{N}\left(0, \sigma_{\tilde{e}}^2\right) \tag{68}$$
$$\tilde{y}_t^\tau = u_{t+1} + \tilde{\eta}_t^\tau \qquad \tilde{\eta}_t^\tau \sim \mathcal{N}\left(0, \sigma_{\tilde{\eta}}^2\right) \tag{69}$$
$$\tilde{\Omega}_t^i = \{x_t, \tilde{p}_t, \tilde{y}_t^\tau\} \cup \tilde{\Omega}_{t-1}^i \tag{70}$$

with $\mathbb{E}\left(u_{t+1}\tilde{e}_t\right) = \mathbb{E}\left(u_{t+1}\tilde{\eta}_t^\tau\right) = 0$. In this formulation, $\tilde{p}_t$ is referred to as a *public signal* and its variance $\sigma_{\tilde{e}}^2$ is used as a measure of the *quality* of the public information, while $\tilde{y}_t^\tau$ is referred to as a *private signal* and its variance $\sigma_{\tilde{\eta}}^2$ is used as a measure its quality.

This section shows that equations (68) and (69) can be derived from equation (66) and that the value of public and private information as condensed by the standard structure can be characterized in terms of the values of σ_p^2, σ_y^2, and σ_ε^2 associated with the model proposed.

First, use (66) to calculate the projection of p_t on u_{t+1}:

$$p_t = a u_{t+1} + e_t \tag{71}$$
$$a = \frac{\mathbb{E}\left(u_{t+1}p_t\right)}{\mathbb{E}\left(u_{t+1}^2\right)} = \frac{\sigma_p^2}{\sigma_p^2 + N\sigma_y^2 + \sigma_\varepsilon^2}$$
$$e_t = p_t - a u_{t+1}$$

This value of a ensures that $\mathbb{E}\left(u_{t+1}e_t\right) = 0$. To write this in the form of (68), divide (71) by a:

[50]Notice that I have already implicitly assumed the observability of x_t, which is most often considered a latent variable in this class of models. See Appendix E discussing latency.

$$a^{-1}p_t = u_{t+1} + a^{-1}e_t$$
$$\tilde{p}_t = u_{t+1} + \tilde{e}_t$$

where $\tilde{p}_t = a^{-1}p_t$ and $\tilde{e}_t = a^{-1}e_t$. Notice by construction that $\mathbb{E}\left(u_{t+1}\tilde{e}_t\right) = 0$ and

$$\begin{aligned}
\tilde{e}_t &= a^{-1}e_t = a^{-1}p_t - u_{t+1} \\
&= \left(\frac{\sigma_p^2 + N\sigma_y^2 + \sigma_\varepsilon^2}{\sigma_p^2}\right) p_t - p_t - \sum_{\tau=1}^{N} y_t^\tau - \varepsilon_{t+1} \\
&= \left(\frac{N\sigma_y^2 + \sigma_\varepsilon^2}{\sigma_p^2}\right) p_t - \sum_{\tau=1}^{N} y_t^\tau - \varepsilon_{t+1}.
\end{aligned}$$

This implies

$$\begin{aligned}
\sigma_{\tilde{e}}^2 &= \frac{\left(N\sigma_y^2 + \sigma_\varepsilon^2\right)^2}{\sigma_p^2} + N\sigma_y^2 + \sigma_\varepsilon^2 \\
&= \frac{\left(N\sigma_y^2 + \sigma_\varepsilon^2\right)\left(N\sigma_y^2 + \sigma_\varepsilon^2 + \sigma_p^2\right)}{\sigma_p^2}
\end{aligned}$$

To build intuition, notice the following limit cases:

$$\lim_{\sigma_p \to 0} \sigma_{\tilde{e}}^2 = \infty$$
$$\lim_{(N\sigma_y + \sigma_\varepsilon) \to 0} \sigma_{\tilde{e}}^2 = 0$$

In words, these limits respectively say that (i) when p_t is always equal to zero, the public signal is worthless; and that (ii) if there are no private signals $\left(y_t^\tau = 0 \text{ with } \sigma_y^2 = 0\right)$ and no purely unpredictable component, the public signal would be perfect, e.g. it would perfectly reveal the future u_{t+1}.

Similarly, we can derive (69). Use (66) to find the projection of y_t^τ on u_{t+1} :

$$\begin{aligned}
y_t^\tau &= bu_{t+1} + \eta_t^\tau \\
b &= \frac{E\left(y_t^\tau u_{t+1}\right)}{E\left(u_{t+1}\right)^2} = \frac{\sigma_y^2}{\sigma_p^2 + N\sigma_y^2 + \sigma_\varepsilon^2} \\
\eta_t^\tau &= y_t^\tau - bu_{t+1}
\end{aligned} \tag{72}$$

where the value of b guarantees that $\mathbb{E}\left(\eta_t^\tau u_{t+1}\right) = 0$. To get this in the form of (69), divide (72) by b :

$$\tilde{y}_t^\tau = u_{t+1} + \tilde{\eta}_t^\tau$$

where $\tilde{y}_t^\tau = b^{-1}y_t^\tau$ and $\tilde{\eta}_t^\tau = b^{-1}\eta_t^\tau$. Analogously:

$$
\begin{aligned}
\tilde{\eta}_t^\tau &= b^{-1} y_t^\tau - u_{t+1} \\
&= \left(\frac{\sigma_p^2 + N\sigma_y^2 + \sigma_\varepsilon^2}{\sigma_y^2} \right) y_t^\tau - p_t - \sum_{\tau=1}^{N} y_t^\tau - \varepsilon_{t+1} \\
&= \left(\frac{\sigma_p^2 + (N-1)\sigma_y^2 + \sigma_\varepsilon^2}{\sigma_y^2} \right) y_t^\tau - p_t - \sum_{\tau' \neq \tau}^{N} y_t^{\tau'} - \varepsilon_{t+1} \\
\sigma_{\tilde{\eta}}^2 &= \left(\frac{\sigma_p^2 + (N-1)\sigma_y^2 + \sigma_\varepsilon^2}{\sigma_y^2} \right)^2 \sigma_y^2 + \sigma_p^2 + (N-1)\sigma_y^2 + \sigma_\varepsilon^2 \\
&= \frac{\left[\sigma_p^2 + (N-1)\sigma_y^2 + \sigma_\varepsilon^2\right] \left[\sigma_p^2 + N\sigma_y^2 + \sigma_\varepsilon^2\right]}{\sigma_y^2}.
\end{aligned}
$$

By taking the limits as for the public signal, we can interpret the following: $\sigma_{\tilde{\eta}}^2 \to \infty$ when $\sigma_y^2 \to 0$. In other words, if the private signal y_t^τ is always zero, the private information is worthless. On the other hand, if there is only one type of signal ($N = 1$) and $\sigma_p^2, \sigma_\varepsilon^2 \to 0$, then $\sigma_{\tilde{\eta}}^2 \to 0$. In this case, the private signal is perfect, or fully revealing ($\equiv \tilde{\sigma}_\eta^2 = 0$).

In summary, the formulation of imperfect/incomplete information that this paper proposed is perfectly consistent with its standard counterpart. The advantages of using this new structure arise from the complete characterization of the data-generating process. Loosely speaking, incomplete information in this form is being generated by something more akin to *asymmetric* information (e.g. *agents don't observe all the info*) than to *imperfect* information (e.g. *agents observe imperfectly a signal of the true data generating process*). It is this aspect of this new structure that allows the researcher to neatly disentangle the effect of the precision of information from the presence of strategic incentives in the preferences of forecasters.

J.1 1-to-1 mapping

This paragraph expands on a remark advanced earlier, that is, the possibility of the two frameworks to be bridged 1-to-1 by means of eliminating ε. The following provides the mapping equations alone, given that the derivations are exactly the same as in the previous paragraph.

The modified information structure is:

$$
\begin{aligned}
x_{t+1} &= \rho x_t + u_{t+1} \\
u_{t+1} &= p_t + \sum_{\tau}^{N} y_t^\tau \\
p_t \sim i.i.d.\left(0, \sigma_p^2\right) &\quad y_t^\tau \sim i.i.d.\left(0, \sigma_y^2\right) \\
\Omega_t^i &= \{x_t, p_t, y_t^\tau\} \cup \Omega_{t-1}^i
\end{aligned}
\tag{73}
$$

By following the same logic as above, the mapping equations are:

$$
\tilde{p}_t = u_{t+1} + \tilde{e}_t \tag{74}
$$

where $\tilde{p}_t = a^{-1} p_t$ and $\tilde{e}_t = a^{-1} e_t$. Again, by construction $\mathbb{E}\left(u_{t+1}\tilde{e}_t\right) = 0$ and

$$
\begin{aligned}
\tilde{e}_t &= a^{-1} e_t = a^{-1} p_t - u_{t+1} \\
&= \left(\frac{N\sigma_y^2}{\sigma_p^2} \right) p_t - \sum_{\tau=1}^{N} y_t^\tau
\end{aligned}
\tag{75}
$$

$$\sigma_{\tilde{e}}^2 = \frac{\left(N\sigma_y^2\right)\left(N\sigma_y^2 + \sigma_p^2\right)}{\sigma_p^2} \tag{76}$$

And analogously

$$\tilde{y}_t^\tau = u_{t+1} + \tilde{\eta}_t^\tau \tag{77}$$

where $\tilde{y}_t^\tau = b^{-1}y_t^\tau$ and $\tilde{\eta}_t^\tau = b^{-1}\eta_t^\tau$. Then:

$$\begin{aligned}
\tilde{\eta}_t^\tau &= b^{-1}y_t^\tau - u_{t+1} \\
&= \left(\frac{\sigma_p^2 + (N-1)\sigma_y^2}{\sigma_y^2}\right) y_t^\tau - p_t - \sum_{\tau' \neq \tau}^{N} y_t^{\tau'} \qquad (78) \\
\sigma_{\tilde{\eta}}^2 &= \frac{\left[\sigma_p^2 + (N-1)\sigma_y^2\right]\left[\sigma_p^2 + N\sigma_y^2\right]}{\sigma_y^2}. \qquad (79)
\end{aligned}$$

The mapping features now a system of two equations in two unknowns which is injective and surjective.

K Chi-Square Test for Over-Identifying Restrictions

The idea behind this note is to illustrate the logic surrounding the estimation of a variance-covariance matrix to be used in the framework of hypotheses testing, and in particular the Chi-square over-identification test.

I expose the case for four ($k = 4$) parameters identified through the use of (at least) five conditions ($n = 5$), these being a mixture of OLS estimated coefficients and variances. By framing the estimated coefficients used for estimation in a Minimum Distance setting, one can ideally calculate the variance-covariance matrix necessary to test hypotheses and make statistical inference.

In total, there are three regressions that summarize the empirical estimates:

$$\begin{aligned} x_{t+h} - \mathbb{F}_t^i\left(x_{t+h}\right) &= \alpha^{p,i} + \beta^p \cdot \left(\mathbb{F}_t^i\left(x_{t+h}\right) - \mathbb{F}_{t-1}^i\left(x_{t+h}\right)\right) + u_t^{p,i} \\ x_{t+h} - \overline{\mathbb{F}}_t x_{t+h} &= \alpha + \beta^c \cdot \left(\overline{\mathbb{F}}_t x_{t+h} - \overline{\mathbb{F}}_{t-1} x_{t+h}\right) + u_t \\ x_{t+h} - \mathbb{F}_t^i(x_{t+h}) &= \alpha_A^{p,i} + \beta_A^p\left(\mathbb{F}_t^i(x_{t+h}) - \mathbb{F}_{t-1}^i(x_{t+h})\right) + \gamma_A^p\left(\mathbb{F}_t^i(x_{t+h}) - \overline{\mathbb{F}}_t(x_{t+h})\right) + u_{A,t}^{p,i} \end{aligned}$$

We focus on the coefficients $\hat{\beta}^c, \hat{\beta}^p, \hat{\beta}_A^p, \hat{\gamma}_A^p$ and on $\hat{\sigma}^2_{u_t^p}$. Rewriting the equations above in deviations from the time (T) average, the five objects of interest are exactly equivalent. In line with the within transformation, define:

$$\widetilde{FE}_{t+h}^i = x_{t+h} - \mathbb{F}_t^i\left(x_{t+h}\right) - \frac{1}{T}\left(\sum_{t=1}^T\left(x_{t+h} - \mathbb{F}_t^i\left(x_{t+h}\right)\right)\right) \tag{80}$$

$$\widetilde{FR}_t^i = \mathbb{F}_t^i\left(x_{t+h}\right) - \mathbb{F}_{t-1}^i\left(x_{t+h}\right) - \frac{1}{T}\left(\sum_{t=1}^T\left(\mathbb{F}_t^i\left(x_{t+h}\right) - \mathbb{F}_{t-1}^i\left(x_{t+h}\right)\right)\right) \tag{81}$$

$$\widetilde{\overline{FE}}_{t+h} = x_{t+h} - \overline{\mathbb{F}}_t x_{t+h} - \frac{1}{T}\left(\sum_{t=1}^T\left(x_{t+h} - \overline{\mathbb{F}}_t x_{t+h}\right)\right) \tag{82}$$

$$\widetilde{\overline{FR}}_t = \overline{\mathbb{F}}_t x_{t+h} - \overline{\mathbb{F}}_{t-1} x_{t+h} - \frac{1}{T}\left(\sum_{t=1}^T\left(\overline{\mathbb{F}}_t x_{t+h} - \overline{\mathbb{F}}_{t-1} x_{t+h}\right)\right) \tag{83}$$

$$\widetilde{DFC}_t^i = \left(\mathbb{F}_t^i(x_{t+h}) - \overline{\mathbb{F}}_t(x_{t+h})\right) - \frac{1}{T}\left(\sum_{t=1}^T\left(\mathbb{F}_t^i(x_{t+h}) - \overline{\mathbb{F}}_t(x_{t+h})\right)\right) \tag{84}$$

Then, we can rewrite the three equations above as follows:

$$\begin{aligned} \widetilde{FE}_{t+h}^i &= \beta^p \cdot \widetilde{FR}_t^i + \tilde{u}_t^{p,i} \\ \widetilde{\overline{FE}}_{t+h} &= \beta^c \cdot \widetilde{\overline{FR}}_t + \tilde{u}_t \\ \widetilde{FE}_{t+h}^i &= \beta_A^p \widetilde{FR}_t^i + \gamma_A^p \widetilde{DFC}_t^i + \tilde{u}_{A,t}^{p,i} \end{aligned}$$

The estimated $\hat{\beta}^c, \hat{\beta}^p, \hat{\beta}_A^p, \hat{\gamma}_A^p, \hat{\sigma}^2_{u_t^p}$ (later I will call these $\hat{\pi}$) are the five data conditions I will focus on in this note. The relative moment conditions (immediately derivable from the F.O.C. of the OLS problem) are:

$$\hat{\beta}^p \qquad \sum_{t=1}^T \sum_{i=1}^{I_t}\left(\hat{\tilde{u}}_t^{p,i} \widetilde{FR}_t^i\right) = 0 \tag{85}$$

$$\hat{\beta}^c \qquad \sum_{t=1}^{T}\left(\hat{\tilde{u}}_t\widetilde{\widetilde{FR}}_t\right) = 0 \tag{86}$$

$$\hat{\beta}^p_A \qquad \sum_{t=1}^{T}\left(\sum_{i=1}^{I_t}\hat{\tilde{u}}^{p,i}_{A,t}\widetilde{FR}^i_t\right) = 0 \tag{87}$$

$$\hat{\gamma}^p_A \qquad \sum_{t=1}^{T}\left(\sum_{i=1}^{I_t}\hat{\tilde{u}}^{p,i}_{A,t}\widetilde{DFC}^i_t\right) = 0 \tag{88}$$

$$\hat{\sigma}^2_{u^p} \qquad \frac{1}{\sum_{t=1}^{T}I_t - 1}\sum_{t=1}^{T}\left(\sum_{i=1}^{I_t}\hat{\tilde{u}}^{p,i^2}_t\right) - \hat{\sigma}^2_{\tilde{u}^p} = 0 \tag{89}$$

By developing $\hat{u}^{p,i}_t$ and $\hat{u}^{p,i}_{A,t}$ along the cross-sectional dimension (I_t) and calling every $\sum_{i=1}^{I_t} x^i_t = x^{Tot}_t$, we just rewrite the above sample moment conditions as:

$$\hat{\beta}^p \qquad \sum_{t=1}^{T}\left[\left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p\widetilde{FR}^{Tot}_t\right)\widetilde{FR}^{Tot}_t\right] = 0 \tag{90}$$

$$\hat{\beta}^c \qquad \sum_{t=1}^{T}\left[\left(\widetilde{\widetilde{FE}}_{t+h} - \hat{\beta}^c\cdot\widetilde{\widetilde{FR}}_t\right)\widetilde{\widetilde{FR}}_t\right] = 0 \tag{91}$$

$$\hat{\beta}^p_A \qquad \sum_{t=1}^{T}\left[\left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p_A\widetilde{FR}^{Tot}_t - \hat{\gamma}^p_A\widetilde{DFC}^{Tot}_t\right)\widetilde{FR}^{Tot}_t\right] = 0 \tag{92}$$

$$\hat{\gamma}^p_A \qquad \sum_{t=1}^{T}\left[\left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p_A\widetilde{FR}^{Tot}_t - \hat{\gamma}^p_A\widetilde{DFC}^{Tot}_t\right)\widetilde{DFC}^{Tot}_t\right] = 0 \tag{93}$$

$$\hat{\sigma}^2_{u^p} \qquad \frac{1}{\sum_{t=1}^{T}I_t - 1}\sum_{t=1}^{T}\left[\left(\hat{\tilde{u}}^{p,2}_t\right)^{Tot}\right] - \hat{\sigma}^2_{\tilde{u}^p} = 0 \tag{94}$$

so to obtain T observations for each sample moment condition.

Define h_t the vector-valued function that stacks the t-evaluated conditions (69) to (73):

$$h_t(\hat{\pi}) = \begin{bmatrix} \left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p\widetilde{FR}^{Tot}_t\right)\widetilde{FR}^{Tot}_t \\ \left(\widetilde{\widetilde{FE}}_{t+h} - \hat{\beta}^c\cdot\widetilde{\widetilde{FR}}_t\right)\widetilde{\widetilde{FR}}_t \\ \left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p_A\widetilde{FR}^{Tot}_t - \hat{\gamma}^p_A\widetilde{DFC}^{Tot}_t\right)\widetilde{FR}^{Tot}_t \\ \left(\widetilde{FE}^{Tot}_{t+h} - \hat{\beta}^p_A\widetilde{FR}^{Tot}_t - \hat{\gamma}^p_A\widetilde{DFC}^{Tot}_t\right)\widetilde{DFC}^{Tot}_t \\ \widetilde{FE}^{2,Tot}_{t+h} + \hat{\beta}^{p,2}\widetilde{FR}^{2,Tot}_t - 2\beta^p(\widetilde{FE}_{t+h}\widetilde{FR}_t)^{Tot} - (I_t - 1)\hat{\sigma}^2_{\tilde{u}^p} \end{bmatrix} \tag{95}$$

Notice that vector-valued function $h_t(\pi)$ is evaluated at $\hat{\pi}$ (where π are the estimands) for each observation t.

Casting the problem as a single GMM estimation problem, we can use the result:

$$\sqrt{T}(\hat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, V_\theta) \tag{96}$$

Assuming that $h_t(\pi_0)$ is serially uncorrelated, we can estimate V_θ as

$$\hat{V}_\theta = (\hat{D}')^{-1}\hat{S}(\hat{D})^{-1} \tag{97}$$

where

$$\hat{D}' = T^{-1} \sum_{t=1}^{T} \frac{\partial h_t(\theta)}{\partial \theta'} \mid_{\theta=\hat{\theta}} = \tag{98}$$

$$= T^{-1} \begin{bmatrix} \sum_{t=1}^{T}(-\widetilde{FR}_t^{Tot} {}^2) & 0 & 0 & 0 & 0 \\ 0 & \sum_{t=1}^{T}(-\widetilde{\overline{FR}}_t^2) & 0 & 0 & 0 \\ 0 & 0 & \sum_{t=1}^{T}(-\widetilde{FR}_t^{Tot} {}^2) & \sum_{t=1}^{T}(-\widetilde{FR}_t^{Tot})(\widetilde{DFC}_t^{Tot}) & 0 \\ 0 & 0 & \sum_{t=1}^{T}(-\widetilde{FR}_t^{Tot})(\widetilde{DFC}_t^{Tot}) & \sum_{t=1}^{T}(-\widetilde{DFC}_t^{Tot} {}^2) & 0 \\ \sum_{t=1}^{T} -2\widetilde{FR}_t^{Tot}(\widetilde{FE}_{t+h}^{Tot} - \hat{\beta}^p \widetilde{FR}_t^{Tot}) & 0 & 0 & 0 & \sum_{t=1}^{T} -2\hat{\sigma}_{\tilde{u}^p} \end{bmatrix}$$

and $\hat{S}$[51]

$$\hat{S} = T^{-1} \sum_{t=1}^{T} h_t(\hat{\pi}) h_t(\hat{\pi})' \tag{99}$$

Then, we can approximate the variance-covariance matrix of $\hat{\theta}$ around θ_0 using $T^{-1}\hat{V}_\theta$. We will use its inverse as weighting matrix for the test statistic of the Chi-square over-identifying restrictions test. The test is as follows:

$$\omega = [\hat{\pi} - \pi(\theta)]' \left(T^{-1}\hat{V}_\theta\right)^{-1} [\hat{\pi} - \pi(\theta)] \sim \chi^2_{n-k} \tag{100}$$

We reject the null if $\omega > \chi^2_{\alpha,n-k}$, where α is the specified level of statistical significance.

K.1 Illustration

For illustration, let's consider the case of a single variable (NGDP) at horizon one. I will display snippets of actual data for purposes of debugging and transparency.

We start by previewing matrix $\hat{H}^1_{NGDP}$, whose rows represent each $h_t(\hat{\pi})'$ from $t = 1$ to $t = T$.

$$\underset{T \times n}{\hat{H}^1_{NGDP}} = \begin{bmatrix} h_1(\hat{\pi})' \\ h_2(\hat{\pi})' \\ \ldots \\ h_T(\hat{\pi})' \end{bmatrix}$$

$$= \underset{204 \times 5}{\begin{bmatrix} 0.0350 & 0.0021 & 0.0109 & 0.0632 & 0.0856 \\ 0.1099 & 0.0024 & 0.0679 & 0.0161 & -0.0567 \\ 0.0246 & 0.0002 & -0.0032 & -0.0228 & -0.1601 \\ -0.0814 & -0.0004 & -0.0446 & -0.0211 & -0.1760 \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix}} \tag{101}$$

Each column in $\hat{H}$ represents the entire series of each sample orthogonality conditions, so that when summed over t they should approximate the population conditions (69) to (73). In fact, they approximate well enough the desired result, with the four OLS coefficients achieving exactly zero, while the residuals' variance estimator getting only close:

$$\sum_{t=1}^{T} \hat{H}^1_{NGDP} = \begin{bmatrix} 0.0000 & 0.0000 & 0.0000 & 0.0000 & -0.0112 \end{bmatrix}$$

[51] The degree of autocorrelation is low, so the use of Newey-West (1987) is second-order. See the Appendix for the autocorrelation coefficients and the autocovariances at lags 1, 2, 3.

$\hat{S}^1_{NGDP}$ is calculated as per (78), and here I report its preview snippet:

$$\underset{n\times n}{\hat{S}^1_{NGDP}} = \underset{5\times 5}{\begin{bmatrix} 0.0709 & 0.0023 & 0.0716 & 0.0010 & 0.0227 \\ 0.0023 & 0.0001 & 0.0024 & -0.0000 & 0.0003 \\ 0.0716 & 0.0024 & 0.0742 & 0.0004 & 0.0224 \\ 0.0010 & -0.0000 & 0.0004 & 0.0014 & -0.0008 \\ 0.0227 & 0.0003 & 0.0224 & -0.0008 & 0.2841 \end{bmatrix}} \tag{102}$$

$\hat{D}^1_{NGDP}$ is calculated as per (77), and here I report its preview snippet:

$$\underset{n\times n}{\hat{D}^1_{NGDP}} = \underset{5\times 5}{\begin{bmatrix} 2.1849 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0027 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 2.1849 & -0.0077 & 0.0000 \\ 0.0000 & 0.0000 & -0.0077 & 0.0118 & 0.0000 \\ -2.0044 & 0.0000 & 0.0000 & 0.0000 & -0.2117 \end{bmatrix}} \tag{103}$$

$\hat{V}^1_{\theta,NGDP}$ is calculated as per (76), and here I report its preview snippet:

$$\underset{n\times n}{\hat{V}^1_{\theta,NGDP}} = \underset{5\times 5}{\begin{bmatrix} 5.2589 & -0.0859 & -0.0284 & 0.3175 & 5.7656 \\ -0.0859 & 11.7498 & 0.4067 & 0.2306 & -0.5187 \\ -0.0284 & 0.4067 & 0.0158 & 0.0603 & -0.0475 \\ 0.3175 & 0.2306 & 0.0603 & 9.9066 & 0.2950 \\ 5.7656 & -0.5187 & -0.0475 & 0.2950 & 6.3386 \end{bmatrix}} \tag{104}$$

$$\underset{n\times n}{\frac{\hat{V}^1_{\theta,NGDP}}{T}} = \underset{5\times 5}{\begin{bmatrix} 0.0258 & -0.0004 & -0.0001 & 0.0016 & 0.0283 \\ -0.0004 & 0.0576 & 0.0020 & 0.0011 & -0.0025 \\ -0.0001 & 0.0020 & 0.0001 & 0.0003 & -0.0002 \\ 0.0016 & 0.0011 & 0.0003 & 0.0486 & 0.0014 \\ 0.0283 & -0.0025 & -0.0002 & 0.0014 & 0.0311 \end{bmatrix}} \tag{105}$$

Finally, equation (79) computes ω^1_{NGDP}:

$$\omega^1_{NGDP} = 0.1452$$

Then, the test can be run using the equations in the body of the paper.

K.2 Autocorrelation of $h_t(\hat{\pi})$

Autocorrelation coefficients for every $h_t(\hat{\pi})$ of one variable (NGDP).

h_t	lag 1	lag 2	lag 3
$h1^1_{NGDP}$	0.0764	-0.0438	0.0609
$h2^1_{NGDP}$	0.0320	0.0033	0.0299
$h3^1_{NGDP}$	0.1134	0.0137	0.0384
$h4^1_{NGDP}$	-0.0188	-0.0584	0.0630
$h5^1_{NGDP}$	0.4436	0.3001	0.2497
$h1^2_{NGDP}$	0.0254	-0.0853	0.1147
$h2^2_{NGDP}$	0.1162	0.0487	0.0553
$h3^2_{NGDP}$	0.0426	0.0207	0.0759
$h4^2_{NGDP}$	0.1641	0.0804	-0.0989
$h5^2_{NGDP}$	0.0582	0.0228	0.0264
$h1^3_{NGDP}$	0.2231	-0.0381	0.1664
$h2^3_{NGDP}$	0.2377	0.1399	0.0490
$h3^3_{NGDP}$	0.1588	0.0654	0.1024
$h4^3_{NGDP}$	0.2427	0.0234	-0.0012
$h5^3_{NGDP}$	0.2114	0.0516	0.0730
⋮	⋮	⋮	⋮
Mean	0.1637	0.0389	0.0336
Median	0.1402	0.0116	0.0127

Table 7: Autocorrelation Coefficients for NGDP

L Limits

L.1 $N \to \infty$ in CG Regression

When taking the limit of $N \to \infty$ before the calculation of β^c, i.e. on consensus forecast error and consensus forecast revisions, we retrieve the expected result of 0 structural covariance ($\beta^c_\infty = 0$).

L.2 $\phi \to -1$ in CG Regression

$$\lim_{\phi \to -1} [\beta^c]$$

$$\lim_{\phi \to -1} \left[\frac{N^2 + N^2\phi - N - N\phi}{(N + \phi(N-1))^2} \right] = 0$$

L.3 $\phi \to 0$ in Augmented Regression

$$\lim_{\phi \to 0} \beta^a = \frac{\sigma_y^4 (N-1)^2}{N^4 \left(\left(\frac{(N-1)^2}{N^2} + \frac{N-1}{N^4} \right) \left(\rho^2 \sigma_\varepsilon^2 + \sigma_y^2 \left(\frac{1}{N^2} + \rho^2 (N-1) \right) + \sigma_p^2 \right) - \frac{\sigma_4^4 (N-1)^2}{N^4} \right)}$$

L.4 $\gamma_A^p = 0$ when $\phi = 0$ in the lagged-consensus specification (eq.(39))

$$\gamma_A^p = \frac{\frac{DQ}{N}\sigma_y^2 \left[\sigma_y^2 \left(\rho^2 Q^2 + \rho^2 (N-1) + D^2\right) + \rho^2 \sigma_\epsilon^2 + \sigma_p^2\right] - \left[\sigma_p^2 + \sigma_y^2 \left(\Lambda Q + \Lambda (N-1) + \frac{D^2}{N}\right)\right] QD\sigma_y^2}{\left[\sigma_y^2 \left(\rho^2 Q^2 + \rho^2 (N-1) + D^2\right) + \rho^2 \sigma_\epsilon^2 + \sigma_p^2\right]\left[\sigma_p^2 + \sigma_\epsilon^2 + \sigma_y^2 \left(N\Lambda + \left(\frac{D}{N}\right)^2\right)\right] - \left[\sigma_p^2 + \sigma_y^2 \left(\Lambda Q + \Lambda(N-1) + \frac{D^2}{N}\right)\right]^2}$$

with $\Lambda = \frac{N^2 - N - \phi(N-1)}{N}$.

Key implication: $\phi = 0 \longrightarrow \gamma_A^p = 0$.

As noted in Appendix D, I have implemented a version of the augmented regression that uses the lagged-consensus to obviate issues of observability. The empirical estimates of such specification are presented in Figure 14. On the theoretical side, the new mapping delivers the above, model-implied expression. The key takeaway from this novel expression is that when $\phi = 0$, $\gamma_A^p = 0$, reconciling our intuition with the internal consistency of the model.[52] However, in the great majority of cases (Figure 14), the result that $\hat{\gamma}_A^p < 0$ holds.

[52] In the background, the forces that were operating towards a negative coefficient – byproduct of the incomplete information environment – are now non-impactful, as there exists a lag between the positive-horizon forecasts and the past consensus estimates. In turn, this is an advantage of the model setup being fundamentally static.

M Gaussianity of information

The assumption of Gaussianity is ubiquitous in most of economics literature that contemplates information acquisition, dissemination or extraction. As reviewed in Appendix J, the standard noisy information structure is attributed to Woodford (2001), but even before then, most studies employing disperse information since Grossman and Stiglitz (1980) have adopted normality as benchmark assumption. The standard structure of information is as follows[53]:

$$\begin{aligned} x_{t+1} &= \rho x_t + u_{t+1} && u_{t+1} \sim \mathcal{N}\left(0, \sigma_u^2\right) \\ \tilde{p}_t &= u_{t+1} + \tilde{e}_t && \tilde{e}_t \sim \mathcal{N}\left(0, \sigma_{\tilde{e}}^2\right) \quad (106) \\ \tilde{y}_t^\tau &= u_{t+1} + \tilde{\eta}_t^\tau && \tilde{\eta}_t^\tau \sim \mathcal{N}\left(0, \sigma_{\tilde{\eta}}^2\right) \quad (107) \\ \tilde{\Omega}_t^i &= \{x_t, \tilde{p}_t, \tilde{y}_t^\tau\} \cup \tilde{\Omega}_{t-1}^i && \quad (108) \end{aligned}$$

In extreme summary, Gaussianity is useful as it is used for[54]:

- **Kalman Filtering**: although not indispensable to generate the Kalman mechanics, Gaussianity is the key assumption guaranteeing the optimality of the algorithm across *any* forecasts functions of observables.
- **Inference**: relatedly in context, but differently in concept, it is worth observing that with Gaussian errors, the Kalman filter's outputs are interpretable as conditional expectations (as opposed to mere linear projections).

Therefore, I ask: *do we have any empirical sense of plausibility of Gaussian signals?* A thorough analysis shows it is an empirically implausible assumption for the vast majority of variables, periods and forecasters. To investigate this matter, I can directly use the SPF's *Probability variables*. These are "auxiliary" variables that are part of each participant's submission allowing to infer an approximation to each forecaster's probability mass function for the variables forecasted at each point in time[55].

As a result, I can build a rich dataset for each of the variables present in the data, and test the assumption of normality. To illustrate, I report the case for variable RGDP (name in the dataset: PRGDP), for which I construct a sample of 4040 pooled observations. I report graphically four of these in Figure 31. Analogous qualitative results obtain for every other variables in the dataset, but I do not include them here for length considerations.

With these newly built datasets, I run two different tests of Gaussianity following the guidance from the statistics research: Doornik and Hansen (2008), D'agostino and Pearson (1973) and Shapiro and Wilk (1965). Table 8 displays a small sample of observations grouped and sorted by forecaster's ID and submission date, together with the average rejection rates. All tests strongly reject normality, with rates of rejection of 90% and 91%.

[53] Here I am maintaining the assumption of x_t being observable, which is most often considered a latent variable in these models. See the section in the Appendix E discussing latency.

[54] For a formal review, see Hamilton (1994) (Ch. 13.4, 13.7, 13.8), Chui, Chen, et al. (2017).

[55] For details, see this link to the SPF.

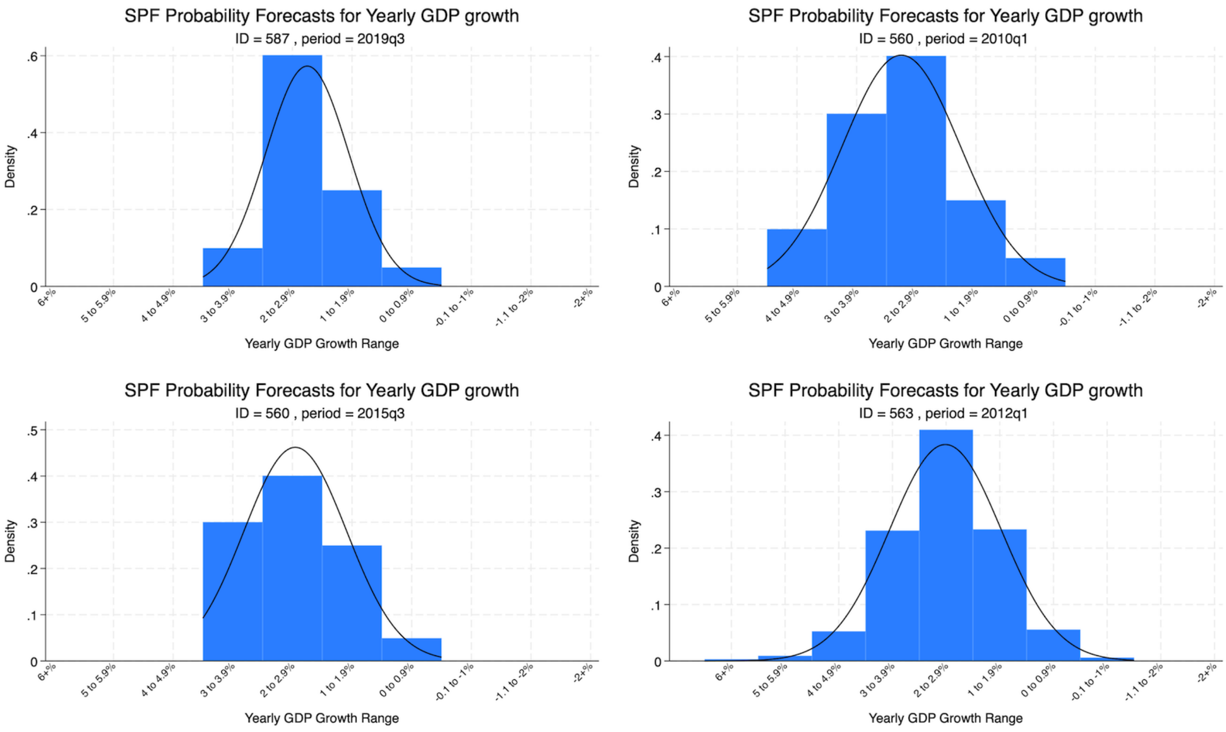

Figure 31: Random sample of probability mass functions.

id	yq	Doornik-Hansen_pval	D'Agostino_pval
20	1992q1	0	0
20	1992q3	1.45e-19	1.88e-13
20	1992q4	.0022126	.0011231
20	1993q1	.0022126	.0011231
20	1993q2	0	0
20	1993q3	0	0
20	1993q4	0	4.52e-37
20	1994q1	0	1.24e-21
20	1994q2	0	1.24e-21
20	1994q4	0	0
20	1995q1	0	1.45e-26
⋮	⋮	⋮	⋮
Avg. Rej. Rate		**.90**	**.91**

Notes: All values are in percentages. Columns 1 and 2 list the sample id and quarter. Columns 3 and 4 report the Doornik-Hansen and D'Agostino test p-values, with rejection rates at the 1% level. Source: SPF, 1968-2019.

Table 8: Sample of Normality Tests with Rejection Rates

N Out-of-Sample (Alternative) Test

Studies by Bianchi et al. (2022) and Eva and Winkler (2023)[56] make an insightful argument that leans against the presence of the significant biases highlighted in the incomplete information literature bloomed after CG (2015). In a nutshell, the argument is that there exists little evidence of the elimination of such biases actually improving the out-of-sample performance of forecasters. In the case of this paper, this would amount to the critique of the substantial deviations I demonstrated. Although there exist meaningful points of agreement between Bianchi et al. (2022) and the strategic perspective this study centers on[57], it could be beneficial to address to the best of possibilities these sensible issues.

However, some technological barriers impede to fully address them. In fact, the most direct way to investigate them would be to run an F-test on the MSEs of two samples of forecasts: (i) the empirical, data-based one; (ii) the model-implied one, obtained by "turning off" the strategic motive of individual forecasters. Conceptually, this is similar to a counterfactual analysis. Unfortunately, such analysis appears unfeasible without stark assumptions, as it is evident from equation (12) in the manuscript. To identify sample (ii), we would necessitate the series for the public and the private information, which amount to $N+1$ unknown series. This same constraint impedes a proper out-of-sample forecast test, as the unobservability of these theoretical objects is unfortunately the norm in the information economics literature.

I implement a second-best alternative that sheds light on these concerns. Ultimately, these hinge upon the quantitative estimation of ϕ, whose average value across variables is ~ 0.5. Although the quantitative estimation of the strategic motive was not the main focus of the paper – as evident from the minimal, parsimonious structure aimed at providing insights more than precise measurement –, a simple experiment to address the comment would be to evaluate forecast performance when "diminishing" the influence of strategic considerations.

In this context, the model estimated a significant incentive to *deviate* from consensus. Then, forecasts *closer* to the consensus should perform better, as they correspond to objective functions that place higher weight on accuracy (i.e., lower ϕ).
Indeed, this turns out to be the case. Using several metrics, Table 9 illustrates this point, comparing individual performance to various aggregate benchmarks. In summary, no individual forecaster significantly beats the equal-weighted (EW) forecast, which therefore sets an upper bound on the accuracy of individual forecasters (in line with the evidence reported in Amodeo, Timmermann and Qu (Working Paper)).

Although not fully representing a counterfactual analysis, this exercise offers some insights to show that strategic incentives are negatively correlated with forecasting accuracy (in a statistically significant way). Columns 2 to 9 represents the rate of rejections with false discover rate not exceeding 5% [58]. Normalized squared-error loss is used when comparing individual forecasts against the benchmark of equal-weighted forecast (EW), cross-sectional average of squared loss

[56] Both these papers focus their criticism on the inability of *lagged* forecasts to predict out-of-sample forecast errors. Notice this *isn't* the specification adopted in the manuscript to discuss strategy; actually, the model predictions for such objects have no predicting power for forecast errors, confirming their findings.

[57] Their conclusions are very much in line with the manuscript's central finding: *"A pervasive finding across all surveys is that respondents place too much weight on the marginal information embedded in their own belief and too little weight on other publicly available information"*.

[58] See Benjamini et al. (2006)

(Average), cross-sectional median of squared loss (Median)[59]. More horizons in Table 10.

Finally, another suggestive indication comes from Figure 9 in the manuscript, which shows that in the absence of strategic incentives the model predicts $\beta^p = 0$, meaning unpredictable forecast errors (i.e., interpretable as efficiency in forecasting).

Table 9: Predictive Accuracy in the SPF: 3Q Horizon

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Number of	Under-perform benchmarks				Out-perform benchmarks			
	forecasters [1]	EW [2]	Average [3]	Median [4]	Rank [5]	EW	Average	Median	Rank
Panel A: SPF (H=1Q)									
CPI	46	1.00	0.02	0.48	0.02	0.00	0.43	0.00	0.00
EMP	32	0.91	0.06	0.59	0.16	0.00	0.66	0.00	0.03
HOUSING	68	0.96	0.04	0.59	0.16	0.00	0.34	0.00	0.03
INDPROD	68	1.00	0.00	0.78	0.06	0.00	0.21	0.00	0.00
PGDP	75	0.97	0.03	0.51	0.07	0.00	0.47	0.00	0.01
RCONSUM	53	1.00	0.02	0.79	0.04	0.00	0.45	0.00	0.00
RFEDGOV	52	1.00	0.00	0.79	0.02	0.00	0.15	0.00	0.02
RGDP	74	1.00	0.04	0.76	0.04	0.00	0.47	0.00	0.03
RNRESIN	52	0.98	0.00	0.73	0.12	0.00	0.37	0.00	0.00
RRESINV	52	0.81	0.02	0.48	0.10	0.00	0.50	0.02	0.04
RSLGOV	52	1.00	0.06	0.71	0.12	0.00	0.40	0.00	0.00
TBILL	50	0.74	0.04	0.72	0.20	0.00	0.52	0.00	0.20
TBOND	46	0.83	0.00	0.65	0.07	0.00	0.35	0.00	0.02
UNEMP	74	0.92	0.03	0.70	0.05	0.00	0.38	0.00	0.08

[1] Only includes individual forecasters with at least 40 observations.
[2] Individual forecasters are compared against the benchmark of the equal-weighted average of individual forecasts at each time period.
[3] The squared loss incurred by individual forecasters are compared against the benchmark loss of the average of individual forecasters' squared error at each time period.
[4] Individual forecasters are compared against the benchmark of the median of individual forecasts at each time period.
[5] The percentile cross-sectional ranks of individual forecasters in terms of squared loss are compared against the benchmark of 0.5.

[59] Specifically, the normalized squared error is adopted to account for unbalanced availability of individual forecasters' forecasts. Let m denote the forecaster; Let $\bar{L}_{mt}$ denote the normalized squared error; L_{mt} denote the squared error; I_{mt} equal to 1 when the forecast by m is available at time t and otherwise 0. The normalized squared error is defined as:

$$\bar{L}_{mt} \equiv \frac{L_{mt}}{\sum_{n=1}^{N} I_{nt} L_{nt} / \sum_{n=1}^{N} I_{nt}}, \quad \text{where} \quad L_{nt} = \left(x_t - \mathbb{F}_{t-h}^{n}(x_t)\right)^2.$$

Of course, the corresponding benchmarks (EW, Average, Median) are also normalized.

Table 10: Predictive Accuracy in the SPF: 0Q and 4Q Horizons

	(1) Number of forecasters [1]	(2) Under-perform benchmarks EW [2]	(3) Average [3]	(4) Median [4]	(5) Rank [5]	(6) Out-perform benchmarks EW	(7) Average	(8) Median	(9) Rank
Panel A: SPF (H=0Q)									
CPI	46	0.87	0.04	0.74	0.30	0.00	0.35	0.04	0.26
EMP	32	0.78	0.03	0.63	0.09	0.00	0.56	0.03	0.28
HOUSING	68	0.93	0.01	0.66	0.09	0.00	0.37	0.00	0.07
INDPROD	67	1.00	0.00	0.78	0.12	0.00	0.30	0.00	0.06
PGDP	75	1.00	0.03	0.67	0.08	0.00	0.51	0.00	0.09
RCONSUM	53	1.00	0.02	0.75	0.04	0.00	0.47	0.00	0.00
RFEDGOV	51	1.00	0.00	0.78	0.08	0.00	0.41	0.00	0.00
RGDP	73	1.00	0.00	0.66	0.08	0.00	0.51	0.00	0.00
RNRESIN	52	1.00	0.02	0.69	0.10	0.00	0.46	0.00	0.04
RRESINV	52	0.96	0.00	0.77	0.06	0.00	0.44	0.00	0.04
RSLGOV	52	1.00	0.06	0.81	0.10	0.00	0.44	0.00	0.00
TBILL	50	0.86	0.04	0.80	0.26	0.02	0.46	0.06	0.24
TBOND	46	0.89	0.02	0.83	0.20	0.00	0.43	0.00	0.13
UNEMP	73	0.97	0.03	0.81	0.10	0.00	0.34	0.00	0.15
Panel B: SPF (H=4Q)									
CPI	44	0.98	0.02	0.39	0.00	0.00	0.39	0.00	0.00
EMP	29	0.93	0.03	0.69	0.14	0.00	0.38	0.00	0.03
HOUSING	55	0.78	0.04	0.38	0.09	0.00	0.29	0.00	0.02
INDPROD	57	0.98	0.02	0.39	0.02	0.00	0.39	0.00	0.00
PGDP	60	0.88	0.02	0.52	0.08	0.00	0.45	0.00	0.05
RCONSUM	49	0.98	0.00	0.73	0.02	0.00	0.39	0.00	0.00
RFEDGOV	46	0.93	0.00	0.54	0.00	0.00	0.24	0.00	0.02
RGDP	60	1.00	0.02	0.58	0.05	0.00	0.33	0.00	0.00
RNRESIN	49	0.84	0.08	0.47	0.12	0.00	0.43	0.00	0.02
RRESINV	49	0.92	0.04	0.78	0.06	0.00	0.31	0.00	0.02
RSLGOV	48	0.94	0.02	0.75	0.10	0.00	0.35	0.00	0.00
TBILL	46	0.46	0.02	0.46	0.13	0.00	0.43	0.00	0.02
TBOND	40	0.70	0.08	0.45	0.13	0.00	0.40	0.00	0.03
UNEMP	61	0.87	0.05	0.61	0.11	0.00	0.28	0.00	0.03

Notes: *see notes from Table 9*

O Underreaction to type-specific public information

Testing the model validates the necessity of the strategic channel for explaining the data. However, an additional intuitive confirmation can come from further empirical exploration of stylized regularities. To address potential criticism of misspecification – or the notion that the strategic force in this model might mask other, concurrent mechanisms – I explore one implication of the model (of many) that was left unused in the estimation, i.e. the evidence of under-reaction to public information. Professional forecasters are typically provided recent official statistics by statistical agencies, including Bureau of Economic Analysis quarterly releases. Surveys like the SPF distribute these latest data points alongside historical values from the previous year or quarters. Forecasters thus base predictions on both updated official releases and recent historical data. Empirical evidence indicates forecasters systematically under-react to public signals and prior beliefs, while overreacting to residual (private) information. This decomposition leverages the availability of consecutive forecasts for the same variable over time.

I employ a similar methodology to Adam et al. (2024) in decomposing forecast revisions into a component predicted by prior beliefs and public information, and a residual capturing private signals. Specifically:

$$\mathbb{F}_t^i\left(x_{t+h}\right)-\mathbb{F}_{t-1}^i\left(x_{t+h}\right)=\alpha_h^{p,i}+\beta_h^p\left(x_t-\mathbb{F}_{t-1}^i(x_t)\right)+\gamma_h^p\mathbb{F}_{t-1}^i\left(x_{t+h}\right)+\epsilon_{t,h} \tag{109}$$

Equation (109) can be read as a simple test of Bayesian forecasting rules, where $\left(x_t-\mathbb{F}_{t-1}^i(x_t)\right)$ represent the arrival of new public but *individual/type*-specific information, and $\mathbb{F}_{t-1}^i\left(x_{t+h}\right)$ the prior belief. Once estimated, the fitted value of regression (109) – which I call $\widehat{\mathbb{FR}}\left(x_{t+h}\right)$ captures adjustments forecasters make based on observable, common information; the residual, $\hat{\epsilon}_t$ represent revisions unexplained by public news and priors, i.e. reflecting private information.

Regressing forecast errors onto these estimated objects allows us to analyze how individuals/-types process public and private innovations:

$$x_{t+h}-\mathbb{F}_t^i\left(x_{t+h}\right)=\zeta_h^{p,i}+\lambda_h^p\widehat{\mathbb{FR}}\left(x_{t+h}\right)+\psi_h^p\hat{\epsilon}_t+\nu_{t,h}^{p,i} \tag{110}$$

Significant coefficients would indicate deviations from rational expectations, such as underreaction or overreaction to public/private signals. This is indeed what we observe across variables and horizons, as highlighted in Figure 32 for horizons 1, 3[60]. Consistently with the rest of the existing evidence and with the model's implications, individual forecasters adjust insufficiently to revisions driven by public information, yet they exhibit an exaggerated response to the residual component. On the contrary, this empirical finding is inconsistent with diagnostic models and their variations, which imply over-reaction to *any* new information.

[60]The rest of available horizons are qualitatively indistinguishable.

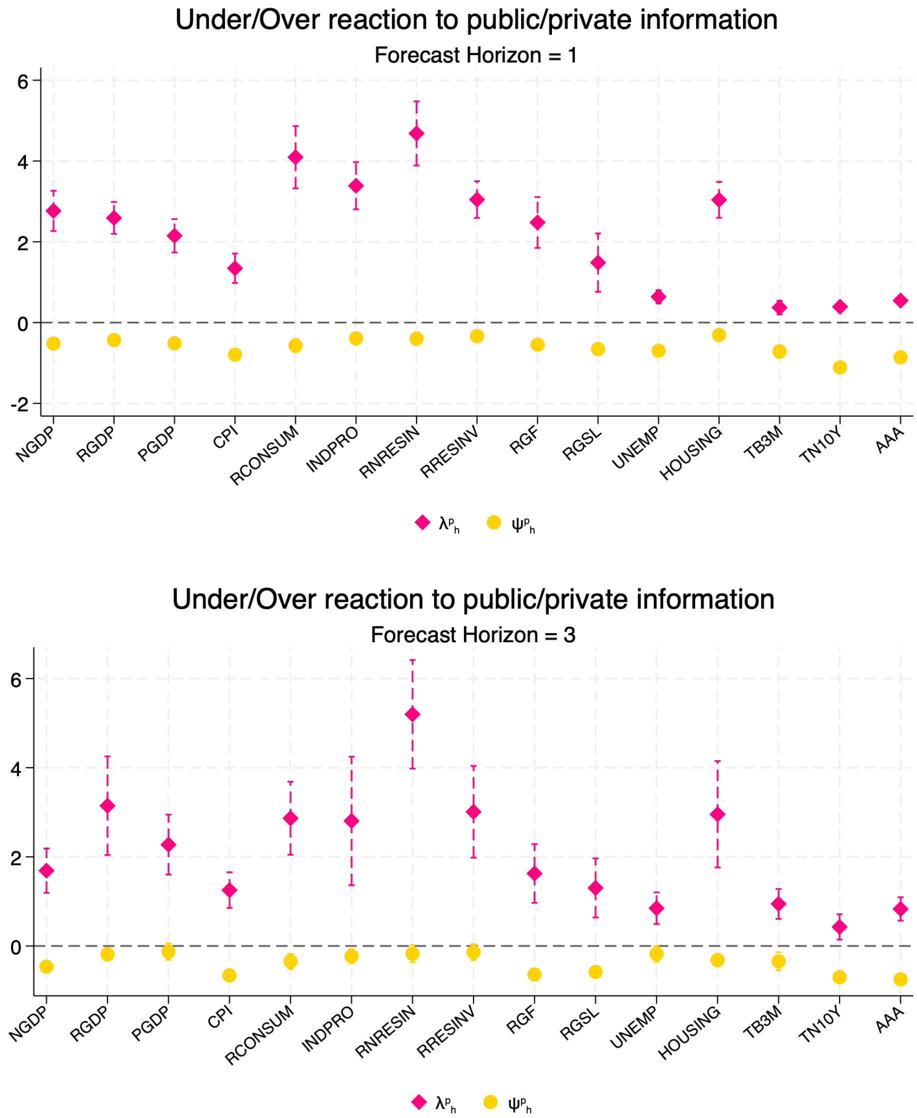

Figure 32: Equation (110)'s λ_h^p and ψ_h^p at horizon 1 and 3, with fixed effects. Confidence intervals at 95% levels.