Salient Forecasters Leading Expectations

Matheus Patrocinio*

June, 2026
(click here for the latest version)

Abstract

Expectations are central to macroeconomics and finance. Standard models assume forecasters operate in isolation, learning only from their own information and from public data, but never from other forecaster. I study a channel where they do learn from each other: when the media reports the inflation forecast of a prominent — *salient* — institution, others can partially infer its idiosyncratic view, so that one institution's private noise can move the whole distribution of expectations. I extend a standard model of expectation formation with what I call *salient signal* and bring it to a panel of Brazilian institution-level inflation forecasts (2010–2021), using a large-language model to identify the days a forecaster is publicized to others and exploiting the differing volatility of market-implied inflation on those days to identify the effect. Forecasters do not react on the day a peer's view becomes news; their expectations then move gradually toward it, reaching a one-for-one response within about three weeks and overshooting thereafter. One institution's idiosyncratic information thus leaves a lasting footprint on aggregate professional expectations.

JEL: D83,D84,E17,E37.

Keywords: Inflation expectations, survey of professional forecasters, information frictions, salience, strategic forecasting, Bayesian learning, media coverage, monetary policy, Brazil.

*Insper matheusmcp@al.insper.edu.br

1 Introduction

Expectations sit at the center of modern macroeconomics and finance, and understanding the process through which professional forecasters form them has become a first-order empirical question. The workhorse models of expectation formation describe forecasters as learning from two sources: their own private information and analysis, and publicly available data such as official releases and policy announcements. Strategic interaction, when it is allowed at all, enters only through forecasters' responses to the aggregate consensus Gemmi and Valchev (2026). In these frameworks, forecasters essentially operate in isolation from one another: there is no channel through which one forecaster directly learns the *view* of another.

Yet such a channel is likely to exist. Some institutions are important enough that the financial press routinely reports their macroeconomic outlooks — big banks, large asset managers, well-known consultancies — and through media coverage and research reports the forecast of one institution becomes visible to all the others. I call such institutions *salient*. What makes their disclosed forecasts economically distinctive is that, unlike a standard public signal such as a data release, they carry the disclosing institution's *idiosyncratic view*: the private information and modelling judgement of one specific forecaster. A salient forecast therefore turns one institution's private signal into an implicit public object — a *salient shock*. If other forecasters learn from it, a single institution's idiosyncratic noise acquires the power to move the entire cross-section of expectations. This possibility is ruled out by construction in standard models, which assume idiosyncratic noises are unobserved and so wash out in the aggregate by the law of large numbers.

This paper asks whether professional forecasters learn directly from salient peers, and how this channel shapes aggregate expectation dynamics. The mechanism is a double-edged sword. On one side, the salient signal works as an *information-aggregation device*: when one institution's view becomes public, others partially incorporate it, compressing the cross-sectional dispersion of beliefs. On the other side, the same channel *amplifies idiosyncratic noise into aggregate noise*: because the disclosed view embeds one institution's private shock, that shock moves all forecasters in the same direction, raising the time-series variance of the consensus. Characterizing this channel — its existence, its magnitude, and its dynamics — requires both a model that makes the mechanism explicit and a research design that can isolate the causal effect of a salient shock in the data.

I proceed in three steps. First, I extend a standard model of strategic expectation formation with a salient signal that arrives *lumpily*: a Bernoulli draw each period determines whether one institution's forecast is disclosed, and, if so, which institution is sampled is governed by exogenous salience weights. The model delivers a closed-form local-projection Jordà (2005) specification in which the object of interest — the impulse response of forecasters' expectations to a salient shock — enters the regression *residual*, because the shock is not directly observed. Second, I identify this response through heteroskedasticity as in Rigobon (2003); Nakamura and Steinsson (2018); Känzig (2021): I use the daily change in Brazilian one-year break-even inflation as an external variable whose variance is higher on disclosure days than on non-disclosure days precisely because of the

salient channel, and the ratio of between-regime differences in covariance and variance recovers the impulse-response function. Third, I take this design to the Brazilian Central Bank's Market Expectations System (the Focus survey), a confidential, institution-level panel of daily inflation forecasts over 2010–2021, and identify disclosure days by applying a large-language-model classifier to Bloomberg Brazilian-economy news headlines.

The estimated response of one-year-ahead inflation expectations to a salient shock is small and statistically indistinguishable from zero on impact: forecasters do not move on the very day a peer's view becomes news, a pattern consistent with the lumpy, infrequent revisions documented by Baley and Turén (2025). The response then builds gradually over the following weeks, crossing one-for-one within roughly three weeks of trading days and reaching about twice the size of the shock by seven weeks before plateauing. The propagation is strongly horizon-dependent: it is largest for short-term (one-year) inflation forecasts and decays by an order of magnitude at the two-, three- and four-year horizons, mirroring the term structure of the heteroskedasticity in the external variable. I read these magnitudes with care — the heteroskedasticity estimator recovers the structural response up to the unknown loading of break-even inflation on the salient shock, so the point estimates are informative about the *sign* and *shape* of the response rather than its exact level.

The central threat to identification is that disclosure days are not random in time: institutions tend to revise, and the press to report, precisely around major inflation releases. I address this through three complementary corrections — controlling for institution-specific IPCA and IPCA-15 release surprises in the local projection, and orthogonalizing the external variable against those surprises. Both corrections leave the estimated impulse-response essentially unchanged in shape and magnitude, which I read as evidence that the estimator recovers the salient-channel response rather than the response to public macro news. A third, more aggressive correction that drops release days from the disclosure sample altogether loses statistical power due to the reduction in the number of disclosure observations.

This paper contributes to four strands of the literature. A first strand is of work on expectation formation, information frictions and biases in survey forecasts. A large body of evidence documents departures from full-information rational expectations — sticky consensus forecasts Coibion and Gorodnichenko (2012, 2015) alongside overreacting individual forecasts Bordalo et al. (2020) — and a complementary line argues these patterns reflect how forecasters *report*, reviving the view that professional forecasters face strategic incentives Laster et al. (1999); Ottaviani and Sørensen (2006): a motive to stand out leads them to overweight idiosyncratic and underweight common information Gemmi and Valchev (2026); overconfidence in private signals produces similar biases Adam et al. (2025); He et al. (2025); and fixed revision costs generate lumpy, consensus-aligned updates Baley and Turén (2025). What unites these frameworks is that forecasters learn only from their own information and from public sources, interacting at most through their response to the aggregate consensus. There is no channel through which one forecaster directly observes, and learns

from, the *view* of another — which is what I add.

Second, the paper extends the literature on salience and the granular origins of aggregate fluctuations Gabaix (2011); Huo et al. (2025) to expectation formation. The mechanism here is granular in the precise sense of Gabaix (2011): one forecaster's idiosyncratic shock, once disclosed, does not wash out in the cross-section but becomes a source of aggregate noise in expectations.

Third, the paper relates to work on news, media, and the formation of expectations York (2023); Binder et al. (2024, 2025), which has studied how media coverage moves *household* expectations. My contribution is to show how media coverage of salient institutions can move *professional* forecasters, the agents whose views feed directly into asset prices and policy.

Finally, the paper contributes methodologically to the literature on high-frequency identification through heteroskedasticity Rigobon and Sack (2004); Nakamura and Steinsson (2018); Känzig (2021), applying the estimator to a new setting: peer-to-peer learning among professional forecasters, where the shock of interest is unobserved and conventional instruments are unavailable.

The remainder of the paper is organized as follows. Section 2 presents the theoretical framework: it introduces the salient signal into a standard strategic expectation-formation model, derives the state-dependent forecasting rule, and characterizes the salient shock and its propagation to aggregate expectations. Section 3 develops the empirical strategy, mapping the model into a local-projection specification, describing the Focus panel and the LLM-based identification of disclosure days, and laying out the heteroskedasticity-based estimator and its identifying assumptions. Section 4 reports the baseline impulse-response function, its behavior across forecast horizons, and the robustness of the estimates to the coincidence of disclosure days with macro releases. The final section concludes.

2 Theoretical framework

In this section I present a model of professional forecasting in which agents learn not only from idiosyncratic and common signals, but also, occasionally, from the publicly disclosed forecasts of *salient* peers. The model is built as a direct extension of a standard expectation formation model with strategic diversification, augmented with a new informational channel that captures one new feature: at unpredictable points in time, the media reports the forecast of a particular institution, and that report enters in the information set of all other forecasters. I refer to this extra signal as the *salient signal* and to the institution-specific noise information it carries as a *salient shock*.

I first describe the target variable process and the two standard signals; then I introduce the salient signal and discuss its institutional content; I then write the forecaster problem, derive its first order condition and discuss its implications and necessary assumptions. The result is a state-dependent solution which decreases the importance of public and private signals whenever the salient signal is revealed, and also opens the possibility that private signal shocks from the salient peers (i.e. the salient shocks) propagate across the other forecasters.

2.1 Target process and the two standard signals

There is a mass of professional forecasters indexed by $i \in \{1, \ldots, N\}$. Each period t, every forecaster reports a prediction $E_t^{(i)}[x_{t+h}]$ of an underlying variable x_{t+h} at horizon h. The variable of interest follows a stationary AR(1) process,

$$x_{t+h} \;=\; \rho\, x_{t-1+h} \;+\; u_t, \qquad u_t \sim \mathcal{N}\big(0, \sigma_u^2\big)\,, \tag{1}$$

with persistence $\rho \in (0,1)$ and innovation variance $\sigma_u^2 \geq 0$.

At the beginning of each period, every forecaster observes three pieces of new information. The first is the **private signal**,

$$s_t^{(i)} \;=\; x_{t+h} \;+\; \eta_t^{(i)}, \qquad \eta_t^{(i)} \overset{\text{iid}}{\sim} \mathcal{N}\big(0, \tau^{-1}\big)\,. \tag{2}$$

The noise $\eta_t^{(i)}$ is independent across agents and over time, with precision τ, and encapsulates the component of i's information set that is uncorrelated with other forecasters' information. We can also think of $\eta_t^{(i)}$ as subjective perception noise or as idiosyncratic modelling choices that produce different point forecasts from the same input data.

The second is the **public signal**,

$$g_t \;=\; x_{t+h} \;+\; \varepsilon_t, \qquad \varepsilon_t \overset{\text{iid}}{\sim} \mathcal{N}\big(0, \nu^{-1}\big)\,, \tag{3}$$

with precision ν and no cross-sectional dimension. This represents the flow of macroeconomic data that is freely available and identical for all agents at time t: high-frequency activity indicators, official statistics, monetary policy releases, and so on.

In the next session I introduce the third, a new signal which is qualitatively different from the last standard ones.

2.2 The salient signal

The novel ingredient of the model is what I call the *salient signal*. Its economic interpretation is that some institutions are sufficiently prominent that the media periodically reports their forecasts. Examples include big banks, large asset managers and well-known consultancies whose macroeconomic outlooks are routinely quoted in the financial press. When such a forecast becomes public, it enters every other forecaster's information set; but, unlike the standard public signal g_t, it carries the disclosing institution's *idiosyncratic view* — that is, the private information of one specific forecaster.

I model this through a two-stage random information arrival.

$1^{st}-$ **Lumpy arrival.** Each period a Berounlli draw defines the current state r_t:

$$r_t \;=\; \begin{cases} R_1 & \text{with probability } q, \\ R_2 & \text{with probability } 1-q. \end{cases} \tag{4}$$

In regime R_1 one institution's forecast is disclosed; in regime R_2 no salient information is released. The realizations $\{r_t\}$ are independent over time.

$2^{nd}-$ **Salience-weighted sampling.** Conditional on $r_t = R_2$, nature draws an institution $k_t \in \{1, \ldots, N\}$ with probability

$$\Pr[k_t = k \mid r_t = R_1] \;=\; w_k, \qquad \sum_{k=1}^{N} w_k = 1, \quad w_k \geq 0. \tag{5}$$

The weights w_k are exogenous and parameterize *salience* — the more prominent a forecaster, the more likely her forecast is to be reported. In the empirical implementation I proxy w_k by newspaper-coverage frequencies, but for now I set this probability as *exogenous* to the model.

The signal. When publication occurs, all agents observe the previous-period reported forecast of the sampled institution, scaled by the persistence of the target:

$$\hat{S}_t^{(k_t)} \;\equiv\; \rho\, \tilde{E}_{t-1}^{(k_t)}[x_{t-1+h}]. \tag{6}$$

Multiplying by ρ shifts the unbiased target from x_{t-1+h} to x_{t+h} via the AR(1) law of motion. The salient signal therefore admits the signal-noise representation

$$\hat{S}_t^{(k_t)} \;=\; x_{t+h} \;+\; \xi_t^{(k_t)}, \qquad \xi_t^{(k_t)} \overset{\text{iid}}{\sim} \mathcal{N}\big(0, \kappa^{-1}\big)\,, \tag{7}$$

where κ is the precision of the salient signal — an *endogenous* object, ultimately pinned down by the steady-state variance of the reported forecast errors, as I show in the closing fixed-point block.

Two features of this construction deserve emphasis.

First, the salient signal is *lumpy* in two senses. It arrives at random intervals (the Bernoulli draw on r_t), and even when it arrives, its content is a single forecaster's view, not an average. This is qualitatively unlike the standard public signal g_t, which is identical across periods in its statistical structure. The lumpy arrival structure is inspired by the fixed-cost view of forecast revisions developed by Baley and Turén (2025), in which information events are sparse rather than continuous.

Second, as we will see next, the salient signal is fundamentally a *private signal made public*. When the model is averaged out, the variance of $\hat{S}_t^{(k_t)}$ traces back to the forecast-error variance of the sampled institution, which in turn is driven primarily by that institution's private signal

noise $\eta^{(k_t)}$. Mechanically, one forecaster's idiosyncratic shock becomes a shock to the entire cross-section of expectations. This is the central economic difference between my framework and the strategic-diversification model of Gemmi and Valchev (2026): in their setting idiosyncratic noises wash out in the consensus, while here a particular institution's $\eta_t^{(k_t)}$ has a non-negligible footprint on aggregate expectation dynamics.

The information set of forecaster i at date t is therefore

$$\mathcal{I}_t^{(i)} \;=\; \{s_t^{(i)}, g_t\} \;\cup\; \{\hat{S}_t^{(k_t)}\}_{r_t=R_1} \;\cup\; \mathcal{I}_{t-1}^{(i)}, \tag{8}$$

augmented with the public observability of r_t and, in regime R_1, of the identity k_t.

2.3 The forecaster's problem

Each forecaster chooses a reported forecast to minimize a quadratic loss that combines accuracy and consensus motives:

$$\min_{E_t^{(i)}} \;\; \mathbb{E}\left[\underbrace{\left(E_t^{(i)} - x_{t+h}\right)^2}_{accuracy} \;-\; \lambda_1 \underbrace{\left(E_t^{(i)} - \bar{E}_t\right)^2}_{strategic} \;\middle|\; \mathcal{I}_t^{(i)} \right], \tag{9}$$

where $\bar{E}_t \equiv \frac{1}{N}\sum_j \tilde{E}_t^{(j)}$ is the consensus reported forecast and $\lambda_1 \in (-\infty, 1)$ parameterizes the strength and sign of strategic incentives. When $\lambda_1 > 0$ the forecaster wishes to "stand out" from the crowd (strategic diversification); when $\lambda_1 < 0$ the forecaster wishes to "herd in" (strategic complementarity). I follow Gemmi and Valchev (2026) in adopting the diversification interpretation as the empirically relevant case, but the algebra goes through for any $\lambda_1 < 1$.

Strategic concerns are present even in pure learning environments because forecasters compete in the marketplace for forecasting expertise — they are not penalized only for being wrong, but also for being indistinguishable from their peers Laster et al. (1999); Ottaviani and Sørensen (2006). Accounting for $\lambda_1 \neq 0$ is essential for reading the data correctly: as Gemmi and Valchev (2026) demonstrate, ignoring strategic incentives could bias the estimated results.

The first-order condition is, taking the derivative of the conditional expectation with respect to $E_t^{(i)}$ and rearranging:

$$E_t^{(i)} \;=\; \frac{1}{1-\lambda_1}\,\mathbb{E}_t^{(i)}\big[x_{t+h}\big] \;-\; \frac{\lambda_1}{1-\lambda_1}\,\mathbb{E}_t^{(i)}\big[\bar{E}_t\big] \tag{10}$$

which is the weighted average between the forecast under rational expectations and the expected average forecasts.

2.4 Restricted Perceptions Equilibrium

Before deriving the optimal forecasting, I must introduce the equilibrium concept this model relies on. A strict rational expectations approach would be intractable here for two reasons. First, the FOC (10) features $\mathbb{E}_t^{(i)}[\bar{E}_t]$, which under full rationality opens an infinite hierarchy of higher-order beliefs — agent i must form a belief about what every other agent believes, and so on. Second, the salient signal $\hat{S}_t^{(k_t)}$ is itself a lagged reported forecast of another agent, so its true noise structure inherits the entire history of cross-sectional forecast errors and regime realizations — a state object whose dimension grows with the sample length. I sidestep both problems by adopting a *Restricted Perceptions Equilibrium* (RPE) in the spirit of Evans and Honkapohja (1993) and Marcet and Nicolini (2003), recently applied to a closely related forecasting environment by Baley and Turén (2025). The idea is that agents act on a *perceived* law of motion for the endogenous objects they cannot fully model, and equilibrium requires this perception to be statistically consistent with what actually realizes. I impose two such restrictions in this paper.

The first concerns the consensus $\bar{E}_t$ that enters the strategic motive. Following Baley and Turén (2025), agents do not solve the infinite hierarchy of higher-order beliefs but instead perceive the consensus as following a simple linear process. Equilibrium then requires that this perceived process be statistically consistent with the actual consensus that arises from aggregating individual reported forecasts.

The second restriction concerns the noise in the salient signal. In truth, $\zeta_t^{(k_t)}$ in (7) inherits the forecast-error dynamics of institution k_t and so has a rich time-series structure. Under the RPE, agents treat it as i.i.d. Gaussian noise with unconditional variance κ^{-1}, and equilibrium requires this perceived variance to equal the actual unconditional variance of $\zeta_t^{(k_t)}$ implied by the model. Without this restriction, the precision matrix in the posterior would depend on the entire history of past regimes and sampled institutions, and the Kalman update would not admit the tractable scalar form I derive next. With it, the consistency of κ reduces to a single scalar fixed-point condition.

2.5 Optimal forecasting

From the first order condition in (10) I now characterize the solution. I start by showing the solution without the strategic motive, i.e., under internally rational expectations. Since the salient signal turns the problem into a state-dependent one, I write the solution in a single regime-indexed expression that nests both states.

Under the RPE, the rational belief $\mathbb{E}_t^{(i)}[x_{t+h}]$ is the Bayesian posterior mean — a Kalman update from the prior $\rho E_{t-1}^{(i)}$ using the available signals:

$$\mathbb{E}_t^{(i)}[x_{t+h}] = \underbrace{\rho E_{t-1}^{(i)}}_{\text{prior}} + \underbrace{G_1^{RE,r_t}}_{\text{gain}} \underbrace{\left(s_t^{(i)} - \rho E_{t-1}^{(i)}\right)}_{\text{innovation}} + G_2^{RE,r_t}\left(g_t - \rho E_{t-1}^{(i)}\right) + G_3^{RE,r_t}\left(\hat{S}_t^{(k_t)} - \rho E_{t-1}^{(i)}\right), \quad (11)$$

with state-dependent Kalman gains $r_t \in \{R_1, R_2\}$:

$$G_1^{RE,r_t} \;=\; \frac{\tau}{D^{RE,r_t}}, \qquad G_2^{RE,r_t} \;=\; \frac{\nu}{D^{RE,r_t}}, \qquad G_3^{RE,r_t} \;=\; \frac{\mathbb{1}\{r_t = R_1\}\cdot\kappa}{D^{RE,r_t}}, \tag{12}$$

where

$$D^{RE,r_t} \;\equiv\; \frac{1}{\Sigma} \;+\; \tau \;+\; \nu \;+\; \mathbb{1}\{r_t = R_1\}\cdot\kappa, \tag{13}$$

and Σ is the steady-state prior variance, endogenouly determined (see Appendix A).

The state-dependence runs entirely through the indicator $\mathbb{1}\{r_t = R_1\}$, which switches the salient signal on and off. In regime R_1 all three signals are active and the posterior weights them by their precisions τ, ν, κ relative to the sum $D^{RE,R_1} = 1/\Sigma + \tau + \nu + \kappa$. In regime R_2 the salient term drops out: $G_3^{RE,R_2} = 0$ and the denominator collapses to $D^{RE,R_2} = 1/\Sigma + \tau + \nu$. Both the gain on the salient signal *and* the gains on the private and public signals are therefore regime-specific — when the salient signal arrives, it competes for posterior weight, shrinking G_1^{RE,R_1} and G_2^{RE,R_1} relative to their counterparts in R_2.

As $\kappa \to 0$ the salient signal becomes pure noise, $G_3^{RE,R_1} \to 0$, and the state-dependence disappears: the agent's posterior collapses to the two-signal case of Gemmi and Valchev (2026). At the other extreme $\kappa \to \infty$, the salient signal is fully informative and $G_3^{RE,R_1} \to 1$ whenever $r_t = R_1$. The empirically relevant case lies between these two limits, and κ itself is pinned down endogenously by the RPE consistency condition introduced earlier and closed at the end of this section.

Equation (11) is the per-regime rational forecast. Reported forecasts depart from it because of strategic incentives, which I now incorporate.

Substituting the rational posterior (11) into the FOC (10) requires characterizing $\mathbb{E}_t^{(i)}[\bar{E}_t]$, the agent's expectation of the average reported forecast. I guess the reported forecast takes the same Kalman form as (11), but with regime-specific reported gains $G_j^{r_t}$ ($j \in \{1,2,3\}$) that differ from their rational counterparts G_j^{RE,r_t}:

$$E_t^{(i)} = \rho E_{t-1}^{(i)} + G_1^{r_t}\big(s_t^{(i)} - \rho E_{t-1}^{(i)}\big) + G_2^{r_t}\big(g_t - \rho E_{t-1}^{(i)}\big) + G_3^{r_t}\big(\hat{S}_t^{(k_t)} - \rho E_{t-1}^{(i)}\big) \tag{14}$$

with $G_3^{R_2} = 0$. Averaging (14) across agents and using the law of large numbers delivers the consensus law of motion, which the RPE requires to coincide with the perceived process. Substituting this back into the FOC and matching coefficients yields the strategic distortion:

$$G_1^{r_t} = \frac{G_1^{RE,r_t}}{D^{r_t}}, \qquad G_2^{r_t} = \frac{(1-\lambda_1)G_2^{RE,r_t}}{D^{r_t}}, \qquad G_3^{r_t} = \frac{(1-\lambda_1)G_3^{RE,r_t}}{D^{r_t}} \tag{15}$$

with $D^{r_t} \equiv 1 - \lambda_1(1 - G_1^{RE,r_t})$.The reported gains are linear transformations of the rational gains, with a regime-dependent distortion factor D^{r_t} that depends only on the rational private gain G_1^{RE,r_t}.

Two implications of (15) are worth flagging.

First, the distortion is not uniform across signals if we focus on the *relative* weights. Because D^{r_t}

depends on G_1^{RE,r_t} but not on G_2^{RE,r_t} or G_3^{RE,r_t} directly, the reported relative weight of the private signal versus the others *does* shift; in fact, under diversification agents systematically overweight the idiosyncratic signal relative to the common ones. This is the central observable implication of Gemmi and Valchev (2026) and is preserved in my model — the new salient signal joins the "common" signal category from this perspective and is similarly underweighted relative to the private signal.

Second, the state-dependence introduces *regime-specific* strategic distortions. Because $G_1^{RE,R_1} \neq G_1^{RE,R_2}$, the factor D^{r_t} takes different values across regimes, and so the strategic wedge between rational and reported gains varies with the information regime. This is a new feature: not only does the salient channel add a new signal, it also alters how strongly agents distort their reported forecasts in different periods.

2.6 The salient shock

Having characterized the reported-forecast rule (14), I can now make precise what is the new piece of information transmitted by the salient signal. Recall that the signal disclosed in regime R_1 is the previous-period reported forecast of the sampled institution, $\hat{S}_t^{(k_t)} = \rho E_{t-1}^{(k_t)}$. Applying (14) one period back to institution k_t:

$$\tilde{E}_{t-1}^{(k_t)} \;=\; \rho E_{t-2}^{(k_t)} \;+\; G_1^{r_{t-1}}\big(s_{t-1}^{(k_t)} - \rho E_{t-2}^{(k_t)}\big) \;+\; G_2^{r_{t-1}}\big(g_{t-1} - \rho E_{t-2}^{(k_t)}\big) \;+\; G_3^{r_{t-1}}\big(\hat{S}_{t-1}^{(k_{t-1})} - \rho E_{t-2}^{(k_t)}\big). \quad (16)$$

Of all the objects appearing on the right-hand side of (16), only *one* is genuinely new information to forecaster i at date t. The public signal g_{t-1}, the salient signal $\hat{S}_{t-1}^{(k_{t-1})}$, and any past public history are already in $\mathcal{I}_{t-1}^{(i)} \subset \mathcal{I}_t^{(i)}$. The prior $\rho E_{t-2}^{(k_t)}$ depends only on k_t's history of beliefs, which i cannot observe directly, but its public-information component is likewise already known. The single piece of k_t-specific information that i does not have is the private signal innovation $s_{t-1}^{(k_t)} - \rho E_{t-2}^{(k_t)}$ — and, through the dependence of $\rho E_{t-2}^{(k_t)}$ on k_t's own past private signals, the accumulated stock of past private shocks specific to k_t.

I define the *salient shock* as this idiosyncratic component of the disclosed forecast:

$$\eta_{t-1}^{(k_t)} \;\equiv\; s_{t-1}^{(k_t)} - x_{t-1+h}, \quad (17)$$

the private signal noise of the sampled institution in the period of its disclosure. By aggregating (14) backwards in time, $\rho E_{t-2}^{(k_t)}$ carries the entire history of past private shocks $\{\eta_{t-1-s}^{(k_t)}\}_{s\geq 1}$ that institution k_t accumulated — I collect these in a residual term $\Delta_{t-2}^{(k_t)}$ capturing k_t's idiosyncratic belief drift relative to the public-information posterior. The publicized forecast then admits the decomposition

$$E_{t-1}^{(k_t)} \;=\; \underbrace{(\text{public-information})}_{\text{known to all forecasters at } t} \;+ G_1^{r_{t-1}}\,\eta_{t-1}^{(k_t)} \;+\; \rho\,\Delta_{t-2}^{(k_t)}. \quad (18)$$

The first term on the right-hand side is already in everyone's information set — it carries no news. The remaining two terms are the new content of the salient signal: the current private innovation $\eta_{t-1}^{(k_t)}$, weighted by the regime-specific gain $G_1^{r_{t-1}}$ that the sampled institution applied at the moment of its disclosure; and the stock of past idiosyncratic shocks $\Delta_{t-2}^{(k_t)}$ that had accumulated in k_t's prior. The salient signal therefore reveals to all forecasters a weighted sum of one institution's accumulated private information.

The salient shock is genuinely idiosyncratic: it is the realization of k_t's own private signal noise plus the residual of her past private signals. Under the standard two-signal information structure, $\eta_{t-1}^{(k)}$ averages out across agents at the consensus level by the law of large numbers and has no aggregate footprint. In my model this is no longer true. When $r_t = R_1$ and $k_t = k$, the entire cross-section of forecasters incorporates k's private information through the social gain $G_3^{R_1}$ in (14). This is the granular channel of aggregate expectation formation that I emphasize throughout the paper, in the spirit of Gabaix (2011).

The transmission of this idiosyncratic information through the salient channel has two opposite effects on aggregate expectation dynamics. On the one hand, the salient signal works as an *information aggregation device*: when one institution's private information becomes public, all forecasters partially incorporate it, reducing the cross-sectional dispersion not only because of the new information but also because the private signal gains are dissolved. On the other hand, the same channel *amplifies idiosyncratic noise* into aggregate noise: the salient shock $\eta_{t-1}^{(k_t)}$ is, by construction, one institution's noise realization, and through $G_3^{R_1}$ it moves the entire cross-section of expectations in the same direction. The time-series variance of the consensus is correspondingly higher than under the standard two-signal information structure, where idiosyncratic shocks wash out by the law of large numbers. The salient channel is therefore a double-edged sword: the very mechanism that aggregates information is also the mechanism that transmits granular noise, and which effect dominates depends on the informativeness of the sampled forecast.

3 Empirical strategy

In this section I bring the model to the data. The theoretical framework of Section 2 delivered a clean impulse-response object: the average response of forecaster i's reported expectation h days ahead to a date $t-1$ *salient shock* $\eta_{t-1}^{(k_t)}$, that is, the private-signal innovation of the institution whose forecast becomes public in regime R_1. The empirical exercise estimates this impulse-response function (IRF) directly from the panel of Brazilian inflation forecasts.

Two challenges shape the strategy. First, the salient shock is *not* directly observed: the Focus survey is confidential, and the econometrician only knows that an institution's view became public on a given day, not the specific private innovation embedded in that disclosure. Second, days on which salient forecasts make the news are not random with respect to other macro news; the very volatility that makes a forecaster newsworthy can coincide with public information releases. Both

challenges are real, and the section is organized around how I deal with them.

I proceed in four steps. Subsection 3.1 describes the data infrastructure: the Brazilian Central Bank's Market Expectations System (Focus survey) as the source of institution-level inflation forecasts, and a LLM-based pipeline applied to Bloomberg Brazilian-economy news headlines to identify the R_1 days on which an institution's forecast is disclosed. Subsection 3.2 maps the model into a local-projection specification Jordà (2005) and reads off the regressors and residual components implied by the theory. Subsection 3.3 addresses the identification of the salient shock: I introduce an external variable z_t constructed from daily changes in Brazilian break-even inflation, argue that z_t is not a valid instrument in the usual sense, and use the heteroskedasticity in z_t across publication and non-publication days to recover the IRF, following Rigobon (2003), Nakamura and Steinsson (2018), and Känzig (2021). Subsection 4.3 discusses the main threat to identification — the coincidence of R_1 days with IPCA release days — and the three robustness tests I implement.

3.1 Data

The Brazilian Expectations survey

The empirical work uses the Market Expectations System maintained by the Brazilian Central Bank (BCB), also known as the Focus survey. The system is one of the most granular panel of professional forecasts available for any major economy. The BCB collects daily institution-level forecasts on a wide menu of variables — inflation indices (IPCA, IGP-M, INPC), GDP, fiscal aggregates, exchange rates, the policy rate and more — at multiple horizons ranging from monthly to several years ahead. Each institution submits its forecasts directly to the BCB, and the average number of active forecasters in a given month is about 95 throughout my sample.

Two institutional features of the survey are essential for the empirical strategy. First, participation and accuracy are incentivized: the BCB publishes monthly and quarterly rankings of the five most accurate forecasters (*Top 5 Focus*), and these rankings carry reputational value in the Brazilian financial industry. Forecasters therefore have an interest in submitting their best view, not a strategic posture designed to game the average. Second — and this is what makes the salient-channel mechanism non-trivial — individual forecasts are *confidential*: only the BCB sees them, and the public release contains only aggregate statistics (mean, median, dispersion). The econometrician sees the panel; the participants do not. Crucially, this means that the only way an institution's specific view reaches its peers is through external channels, of which media coverage is the most natural one. The salient channel of Section 2 is therefore a description of this information environment.

My sample period is 2010–2021. I restrict attention to forecasts of one-year-ahead IPCA inflation, the main consumer price index targeted by the BCB. This horizon is the one most regularly reported in the financial press. Since the survey provides end-of-year forecast horizons, I make linear interpolation of forecasts in order to have exact one-year-ahead expectations.

Identifying R_1 days from Bloomberg news

To bring the publication regime R_1 from definition in Equation (4) to the data, I need to identify the dates on which one institution's inflation view was disclosed publicly via the media. My source is the Bloomberg *Brazilian Economy News* (ticker *BZECO*) feed, which covers the top thirty most-read news items per day over the period 2010–2021, for a total of 97,996 headlines.

A traditional approach to this kind of text-classification problem would be a keyword search, but the false-positive rate is high: many inflation-related headlines are not about institutional forecast revisions, and many forecast-revision headlines do not contain obvious keywords. I therefore use a small large-language model (Mistral Small 3.2) as a classifier. The model receives the headline text and a prompt asking it to decide whether the headline describes a financial institution explicitly revising its Brazilian inflation forecast. The prompt is constructed with both positive and negative examples. The full prompt is reported in Appendix B.1.

The pipeline identifies 439 qualifying news distributed across 328 days in which 258 are business days, out of 2,908 total trading days in the sample. These 258 days are my empirical analogue of R_1: dates on which an institution-level inflation forecast became public via the press. The remaining 2,650 days are my R_2 sample.

Figure 7 in Appendix B reports the distribution of news items across institutions; eight institutions account for over half of the coverage, with Banco Itaú (64 news), Ibiuna Investimentos (56), Barclays (53), Credit Suisse (28), Santander (27), Bradesco (22), MCM Consultores (14), and BNP Paribas (12) leading the list.

3.2 The Local Projections Specification

The empirical question I want to answer is how the cross-section of professional inflation forecasts responds to a salient shock — that is, to the realization of the private-signal innovation $\eta_{t-1}^{(k_t)}$ of the institution whose forecast is sampled and disclosed in period t. Section 2 established that this shock is the only piece of genuinely new information embedded in the salient signal, and that, unlike under the two-signal benchmark, its footprint does not average out across forecasters: through the salient gain $G_3^{R_1}$, it moves the entire distribution of reported expectations.

The natural object that summarizes this propagation is the impulse-response function (IRF) of forecaster i's reported expectation at horizon h to a salient shock realized at $t-1$:

$$\mathrm{IRF}_h^{(i,k_t)} \;\equiv\; \frac{\partial E_{t+h}^{(i)}}{\partial \eta_{t-1}^{(k_t)}}. \tag{19}$$

Two features of (19) call for comment before I can write an econometric specification for it.

First, the IRF depends on a *stochastic* path. After date $t-1$, each subsequent period $\ell \in \{t, t+1, \ldots, t+h\}$ draws a fresh regime $r_\ell \in \{R_1, R_2\}$ from the Bernoulli structure of Section 2, and in R_1 periods a fresh sampled institution k_ℓ is drawn from the salience distribution $\{w_k\}$. The

dynamic gains and persistence parameters that propagate the shock will depend on the realized future regime sequence. The path of forecaster i's posterior between $t-1$ and $t+h$ therefore varies with the realized future history $(r_t, k_t, r_{t+1}, k_{t+1}, \ldots)$, and so does (19).

Second, the same is true of the identity of the institution that generated the shock. The IRF in (19) is conditional on $k_t = k$; under the symmetry of the model, the magnitude of the response is the same across institutions k when one normalizes by the salience weight, but the conditional IRF still depends on which institution was sampled at t.

What an empirical exercise can recover is the *average* of the IRF over both sources of stochasticity: averaging over the realized future regimes $(r_t, \ldots, r_{t+h})$ and over which institution was sampled at t. This is the unconditional IRF $\mathbb{E}\left[\frac{\partial E^{(i)}_{t+h}}{\partial \eta^{(k_t)}_{t-1}}\right]$ where the expectation integrates jointly over the future regime path and the sampling distribution. It is what a local-projection regression Jordà (2005) estimates — a regression that does not condition on the realized future history but instead averages over it sample-period by sample-period. In what follows I show how the analytical model maps into a closed-form LP specification.

The starting point is the reported-forecast rule (14) of Section 2, which expresses $E^{(i)}_t$ as a Kalman update from the prior with the three signals available in the current regime. The rule is regime-dependent: the salient signal $\hat{S}^{(k_t)}_t = \rho E^{(k_t)}_{t-1}$ enters with gain $G^{R_1}_3$ in R_1 and is absent in R_2. Taking the expectation over the publication state r_t, the regime-dependent gains collapse to their regime-averaged counterparts and the rule becomes

$$\mathbb{E}\big[E^{(i)}_t \,\big|\, s^{(i)}_t, g_t\big] \;=\; \rho E^{(i)}_{t-1} \;+\; \bar{G}_1\big(s^{(i)}_t - \rho E^{(i)}_{t-1}\big) \;+\; \bar{G}_2\big(g_t - \rho E^{(i)}_{t-1}\big) \;+\; \bar{G}_3\big(\rho \mathbb{E}[E^{(k_t)}_{t-1}] - \rho E^{(i)}_{t-1}\big), \quad (20)$$

where the regime-averaged gains are defined as

$$\bar{G}_p \;\equiv\; q\, G^{R_1}_p + (1-q)\, G^{R_2}_p, \qquad p \in \{1,2,3\}, \quad (21)$$

with $\bar{G}_3 = q\, G^{R_1}_3$ since $G^{R_2}_3 = 0$, and the expectation in the last term integrates over the salience-weighted sampling distribution of the disclosed institution: $\mathbb{E}[E^{(k_t)}_{t-1}] \equiv \sum_k w_k\, E^{(k)}_{t-1}$.

The trouble with (20) as it stands is that the expected salient signal $\mathbb{E}[E^{(k_t)}_{t-1}]$ is a single object that bundles together consensus, the private innovation of the sampled institution, and the accumulated idiosyncratic drift inherited from previous publications. The empirical exercise needs these three pieces separately — the consensus piece becomes a regressor, while the latter two pieces are the salient-shock components that I want to isolate inside the residual. I therefore open $\mathbb{E}[E^{(k_t)}_{t-1}]$ using the same decomposition as in equation (18) of Section 2, integrated over the salience distribution of k_t:

$$\mathbb{E}\big[E^{(k_t)}_{t-1}\big] \;=\; \underbrace{\bar{E}_{t-1}}_{\text{consensus}} \;+\; \underbrace{\bar{G}_1\, \mathbb{E}[\eta^{(k_t)}_{t-1}]}_{\text{salient shock}} \;+\; \underbrace{\rho\, \bar{\psi}\, \mathbb{E}[\Delta^{(k_t)}_{t-2}]}_{\text{accumulated drift}} . \quad (22)$$

The first term is the equally-weighted consensus of all institutions' lagged forecasts and is directly observable in the Focus data. The second term is the private-signal innovation of the sampled

institution scaled by the regime-averaged private gain $\bar{G}_1$ — this is the *salient shock* of Section 2. The third term is the accumulated belief drift $\Delta_{t-2}^{(k_t)} \equiv E_{t-2}^{(k_t)} - \bar{E}_{t-2}$ of the sampled institution, scaled by the persistence factor $\rho\bar{\psi}$. The decomposition uses two regime-averaged eigenvalues that will be useful below,

$$\bar{\psi} \;\equiv\; 1 - (\bar{G}_1 + \bar{G}_2 + \bar{G}_3), \qquad \bar{\phi} \;\equiv\; \bar{\psi} + \bar{G}_3, \tag{23}$$

which govern, respectively, the average decay of an institution's own idiosyncratic deviation from the consensus and the average decay of the consensus itself under the salient channel.

Substituting (22) into (20) and collecting terms by observability yields the explicit one-period expected propagation:

$$\begin{aligned}\mathbb{E}\big[E_t^{(i)} \,\big|\, s_t^{(i)}, g_t\big] = \rho\bar{\psi}\, E_{t-1}^{(i)} + \bar{G}_1\, s_t^{(i)} + \bar{G}_2\, g_t + \rho(\bar{\phi} - \bar{\psi})\, \bar{E}_{t-1} \\ + \rho\, \bar{G}_1(\bar{\phi} - \bar{\psi})\, \mathbb{E}[\eta_{t-1}^{(k_t)}] + \rho^2\bar{\psi}(\bar{\phi} - \bar{\psi})\, \mathbb{E}[\Delta_{t-2}^{(k_t)}].\end{aligned} \tag{24}$$

The first four terms collect the observable regressors and signals that the LP will absorb on the right-hand side; the last two terms collect the unobservable salient-shock components that will constitute the residual. The same logic applies, with appropriate modifications, when projecting the forecast forward to horizon h: the structure of (24) is preserved, with the coefficients on each piece becoming the corresponding h-step powers of ρ and the eigenvalues $\bar{\phi}, \bar{\psi}$.

Augmenting the conditional expectation with institution and month fixed effects to absorb level heterogeneity, the **local-projection regression** that follows from the model is

$$E_t^{(i)}[\pi_{t+h}] \;=\; \alpha_h^{(i)} \;+\; \phi_h^m \;+\; \beta_{1,h}\, E_{t-1}^{(i)} \;+\; \beta_{2,h}\, \bar{E}_{t-1} \;+\; u_{t+h}^{(i)}. \tag{25}$$

The two regressors are observable: $E_{t-1}^{(i)}$ is forecaster i's own lagged forecast, and $\bar{E}_{t-1}$ is the equally-weighted consensus, both directly available in the Focus data. The fixed effects $\alpha_h^{(i)}$ and ϕ_h^m absorb individual and monthly fixed effects respectively.

The interest of the LP is actually in its *residual*. Reading $u_{t+h}^{(i)}$ through the model and collecting the structural components yields the decomposition

$$u_{t+h}^{(i)} \;=\; \omega_{1,h}\, s_t^{(i)} \;+\; \omega_{2,h}\, g_t \;+\; \underbrace{\omega_{3,h}\, \eta_{t-1}^{(k_t)}}_{\text{salient shock}} \;+ \omega_{4,h}\, \Delta_{t-2}^{(k_t)} \;+\; \xi_{t+h}^{(i)}. \tag{26}$$

The five terms collected on the right-hand side are: (i) the forecaster's own current-period private signal $s_t^{(i)}$, weighted by its propagation through the model to horizon h; (ii) the current public signal g_t, similarly weighted; (iii) the salient shock $\eta_{t-1}^{(k_t)}$ that was disclosed at t, weighted by the IRF coefficient $\omega_{3,h}$; (iv) the accumulated idiosyncratic belief drift $\Delta_{t-2}^{(k_t)}$ of the sampled institution at the time of disclosure, as defined in equation (18) of Section 2; and (v) the orthogonal noise term $\xi_{t+h}^{(i)}$ collecting future innovations realized between t and $t+h$ (target innovations $u_{t+\ell}$, future public shocks $\varepsilon_{t+\ell}$, future private shocks $\eta_{t+\ell}^{(i)}$, and future salient draws).

The coefficient of interest is $\omega_{3,h}$ — the loading of the LP residual on the salient shock

$$\omega_{3,h} \;=\; \mathbb{E}\left[\frac{\partial E^{(i)}_{t+h}}{\partial \eta^{(k_t)}_{t-1}}\right]. \tag{27}$$

If the salient shock $\eta^{(k_t)}_{t-1}$ were observable, an identification strategy would be immediate. But $\eta^{(k_t)}_{t-1}$ is the private-signal noise of an institution that the econometrician cannot recover from the public record: even on R_1 days, the *magnitude* of the institution's private innovation is unknown, only the fact that one was disclosed. The remaining components of the residual — the current-period private signal $s^{(i)}_t$, the current public signal g_t, the lagged drift $\Delta^{(k_t)}_{t-2}$, and the future innovations in $\xi^{(i)}_{t+h}$ — are likewise unobserved.

The strategy I adopt is to find an external variable that loads on $\eta^{(k_t)}_{t-1}$ and to exploit the heteroskedasticity of that variable across the publication and non-publication regimes. The next subsection develops this argument.

3.3 Heteroskedasticity-based identification

I use as external variable z_t the daily variation in the one-year ahead break-even inflation rate implied by Brazilian Treasury bond prices, computed by ANBIMA,

$$z_t \;\equiv\; \Delta\pi^{\mathrm{BE},1y}_t. \tag{28}$$

The one-year break-even is the difference between nominal and inflation-indexed Treasury yields at the one-year horizon, and moves daily with market-implied inflation expectations. I expect break-even inflation responds to events such as macro releases (IPCA, IPCA-15, exchange rate, commodity prices), policy news (COPOM decisions, fiscal announcements), and also to the publication of forecasts by salient institutions. The model interpretation of such variable is the linear decomposition

$$z_t \;=\; a\,\eta^{(k_t)}_{t-1} \;+\; b\,\varepsilon_t \;+\; \nu_t, \tag{29}$$

where the first term captures the market's reaction to a salient shock revealed at t (with loading a that depends on how strongly break-even inflation reflects professional-forecast revisions), the second is the market's reaction to other public macro news with loading b, and ν_t collects measurement errors.

A naive instrumental-variable strategy — treating z_t as an instrument for the salient shock in (26) — fails on its face: equation (29) makes plain that z_t is correlated with the public-signal component ε_t of the residual, and the exclusion restriction is violated. What I exploit instead is the differential variance of z_t across R_1 and R_2: under the structural decomposition, the salient-shock component of z_t has positive variance in R_1 and zero variance in R_2, while the public-shock and orthogonal components have variances that are, by assumption, the same across regimes. This is

the canonical setup of identification through heteroskedasticity Rigobon (2003).

The identification rests on three conditions on the variances of the components of the LP residual and the proxy across the two regimes. Throughout, let $\sigma^2_{x,r}$ denote $\mathrm{Var}(x \mid r_t = r)$.

(A1) **Differential variance of the salient channel:** $\sigma^2_{\eta,R_1} \neq \sigma^2_{\eta,R_2}$. In the model this holds analytically: in R_2 no institution is sampled, so the salient shock $\eta^{(k_t)}_{t-1}$ has zero variance contribution to z_t, while in R_1 it has variance $\tau^{-1} > 0$.

(A2) **Common variance of public macro shocks:** $\sigma^2_{\varepsilon,R_1} = \sigma^2_{\varepsilon,R_2}$. The public-information component of z_t has the same variance on R_1 and R_2 days. This is a substantive restriction — it says that the days on which an institution's forecast becomes news are not systematically days of more (or less) macro-news volatility than days on which no institution makes the news.

Conditions (A1)–(A2) together guarantee that the differential variance of z_t across R_1 and R_2 isolates the salient channel, and that the differential covariance of z_t with $u^{(i)}_{t+h}$ across regimes isolates the effect of the salient shock on future forecasts.

Following Rigobon (2003); Nakamura and Steinsson (2018); Känzig (2021), the estimator of the IRF is the ratio of between-regime differences in covariance and variance:

$$\widehat{\omega}_{3,h} = \frac{\widehat{\mathrm{cov}}_{R_1}(z_t,\, \bar{u}_{t+h}) - \widehat{\mathrm{cov}}_{R_2}(z_t,\, \bar{u}_{t+h})}{\widehat{\mathrm{var}}_{R_1}(z_t) - \widehat{\mathrm{var}}_{R_2}(z_t)} \approx \mathbb{E}\left[\frac{\partial E^{(i)}_{t+h}}{\partial \eta^{(k_t)}_{t-1}}\right], \tag{30}$$

where $\bar{u}_{t+h} \equiv N^{-1}\sum_i \hat{u}^{(i)}_{t+h}$ is the cross-sectional average of the LP residuals from (25).[1] Under (A1)–(A2), the numerator is a consistent estimate of $\omega_{3,h} \cdot a \cdot [\sigma^2_{\eta,R_1} - \sigma^2_{\eta,R_2}]$ and the denominator is a consistent estimate of $a^2 \cdot [\sigma^2_{\eta,R_1} - \sigma^2_{\eta,R_2}]$, so their ratio is $\omega_{3,h}/a$. The proxy loading a rescales the magnitude of the estimated IRF but not its sign or shape; for the impulse-response interpretation it is convenient to normalize $a = 1$, in which case $\widehat{\omega}_{3,h}$ recovers the average causal effect on day $t+h$ of a unit salient-shock innovation at $t-1$.

A precondition for the estimator (30) is that the denominator be statistically different from zero — that is, that the proxy z_t actually exhibits more variation in R_1 than in R_2. I test this directly. Table 1 reports the variance of z_t in each regime, the variance ratio $\sigma^2_{R_1}/\sigma^2_{R_2}$, the Brown–Forsythe F-statistic and p-value, and the number of observations in each regime, for the one-year break-even and for break-evens at the 2-, 3-, and 4-year horizons.

Three patterns emerge from Table 1. First, the one-year break-even is significantly more volatile on R_1 days than on R_2 days: the variance ratio is 1.40 with a p-value below 0.001. The Brown–Forsythe test, which is robust to non-normality, is firmly inside the rejection region. Second, the strength of the heteroskedasticity declines monotonically with the horizon of the break-even, with

[1] Pooling the residual across i before forming the covariance reduces idiosyncratic noise without affecting the population identification, since $\omega_{3,h}$ is the same across institutions.

Table 1. Heteroskedasticity test for the external variable z_t.

	σ_{R_1} (bps)	σ_{R_2} (bps)	$\sigma^2_{R_1}/\sigma^2_{R_2}$	F-stat	p-value	N_{R_1}	N_{R_2}
$\Delta\pi_t^{\text{BE},1y}$	9.07	7.65	1.40	11.07	0.00	258	2650
$\Delta\pi_t^{\text{BE},2y}$	7.15	6.26	1.30	8.27	0.00	258	2650
$\Delta\pi_t^{\text{BE},3y}$	7.02	6.49	1.17	4.47	0.03	258	2650
$\Delta\pi_t^{\text{BE},4y}$	6.96	6.82	1.04	1.58	0.21	258	2650

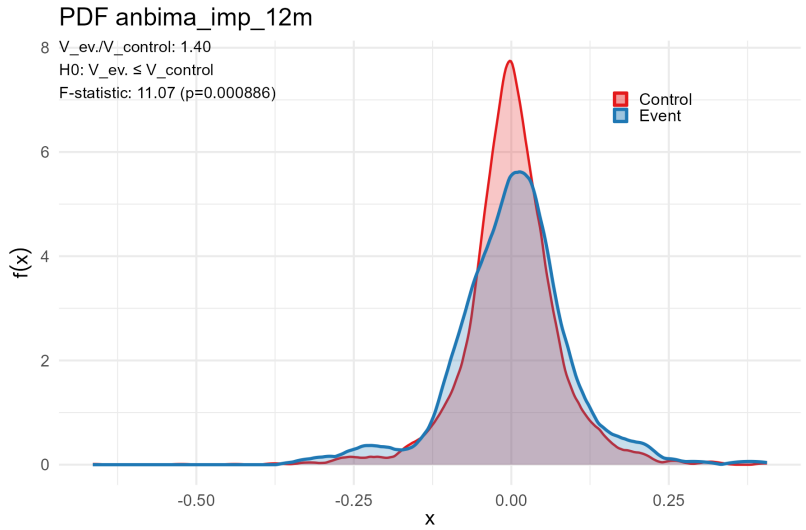


Figure 1. Empirical PDF for $\Delta\pi_t^{BE,1y}$ estimated using Epanechikov kernel

the four-year break even no longer statistically distinguishable from one. This pattern is consistent with the salient channel having most of its effect at shorter-term horizons, where institutional inflation forecasts carry the strongest news content for market pricing. Another way to visualize the heteroskedasticity is through the probability density function in Figure 1 which displays a flatter distribution of daily variation in days with events.I therefore proceed with the one-year break-even as the external variable z_t.

4 Results

This section reports the main empirical findings. I begin with the baseline impulse-response function (IRF) of forecasters' one-year-ahead inflation expectations to a salient shock, estimated from the heteroskedasticity-based local-projection procedure developed in Section 3. I then extend the exercise to longer forecast horizons — 2-, 3- and 4-year ahead — and show that the salient-channel response decays with the horizon of the forecast, as the model would predict. Finally, in Subsection 4.3 I address the most direct threat to the identification — the systematic coincidence of R_1 days with IPCA release days — through three complementary tests, and show that the baseline

IRF is essentially unchanged.

4.1 Effect of salient shocks on inflation expectations

Figure 2 reports the baseline estimate of $\omega_{3,h}$ for the one-year-ahead inflation forecast, as a function of the horizon h measured in trading days. Since I use end-of-year projections, I have to make a linear interpolation between current year and next year forecasts to access what would be the projected forecast for exactly one year ahead. The point estimate is the heteroskedasticity ratio of equation (30), computed on the cross-sectional average of the LP residuals from (25); the shaded bands report the 68% and 90% confidence intervals.

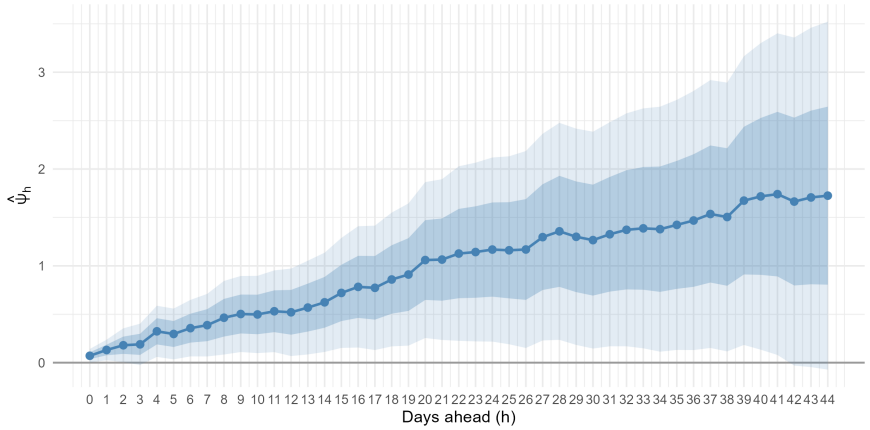


Figure 2. Estimated response of 1-year-ahead inflation expectations to a salient shock, by horizon h (in trading days). Shaded areas report 68% and 90% bootstrap confidence bands.

The response is small and statistically indistinguishable from zero on impact. At $h = 0$, $\widehat{\omega}_{3,0} \approx 0.1$ and the confidence band comfortably includes zero. Forecasters do not adjust their reported views on the very day a peer's forecast becomes news — a pattern consistent with the lumpy-revision behaviour documented by Baley and Turén (2025), in which forecasters do not revise continuously but in discrete bursts. The response then builds gradually over subsequent weeks as more forecasters update their views.

Also, the IRF grows monotonically and crosses one within approximately three weeks of trading days. By $h \approx 21$ the average forecaster's expectation has moved one-for-one with the salient shock: the typical institution effectively incorporates a unit shock to a salient peer's view within fifteen trading days of its disclosure. By $h \approx 35$ (about seven weeks) the response roughly doubles, reaching $\widehat{\omega}_{3,h} \approx 2$, and then plateaus around 1.8–1.9 at the longest horizons I can estimate, $h \approx 44$.

Finally, while the medium-horizon magnitudes are striking at face value, their structural interpretation requires care. As noted in Section 3.3, the heteroskedasticity-based estimator does not recover the structural IRF coefficient $\omega_{3,h}$ directly, but rather the ratio $\omega_{3,h}/a$, where a is the loading of the external variable z_t on the salient shock in the decomposition (29). The point es-

timates in Figure 2 are therefore informative about the *sign* and *shape* of the response, but their absolute level conflates the structural response with the rescaling factor a. If $a < 1$ — which is the empirically plausible case, since the one-year break-even is unlikely to absorb the full content of a salient innovation one-for-one — the estimator mechanically inflates the structural $\omega_{3,h}$, and the apparent crossing of unity at three weeks or of two at seven weeks does not by itself constitute evidence of overreaction in the structural sense. What the baseline IRF establishes robustly is that the salient-channel response is positive, statistically significant, and builds gradually over several weeks — a qualitative pattern consistent with the salient-signal mechanism of Section 2, and in particular with the strategic-diversification reading of Gemmi and Valchev (2026) operating through the new channel. Disentangling the scaling factor a from the structural response is a natural next step, and I leave it for future work.

4.2 Effects across forecast horizons

A natural follow-up is whether the salient channel propagates symmetrically across forecast horizons. The Focus survey reports institution-level inflation forecasts at multiple horizons, and the same heteroskedasticity-based LP can be applied to the 2-, 3- and 4-year-ahead forecasts holding the external variable z_t (the one-year break-even) and the regime definition fixed. The exercise tests how far down the term structure of expectations the salient-channel impulse travels.

Figure 3 reports the four IRFs side by side. The pattern is clear: the response is largest, in absolute magnitude, at the one-year horizon and declines sharply as the forecast horizon lengthens. At the 2-, 3- and 4-year horizons the IRF eventually reaches a peak of roughly 0.2 by trading-day forty, but the on-impact response and the early dynamics are essentially flat, and the long-horizon point estimate is an order of magnitude smaller than its one-year counterpart.

This pattern is consistent with the heteroskedasticity test of Table 1: the proxy z_t exhibits the strongest between-regime heteroskedasticity at the one-year horizon, and the longer-horizon break-evens are less informative about salient-channel shocks — the four-year break-even, in particular, is not statistically distinguishable across regimes. The interpretation is that institutional inflation forecasts disclosed in the press carry their most pricing-relevant content for short-term expectations, where the universe of professional forecasters appears most responsive, while long-horizon expectations are anchored by more persistent factors (target levels, regime credibility) that react little to the disclosure of an individual institution's view.

4.3 Robustness to macro releases

The most direct threat to the identifying assumption (A2) — common variance of public macro shocks across regimes — is that R_1 days are not randomly distributed in time. Of the 258 R_1 days in my sample, 58 coincide with monthly IPCA inflation release and 49 coincide with monthly IPCA-15 release, which is a preliminary measure release for the inflation of that specific month.

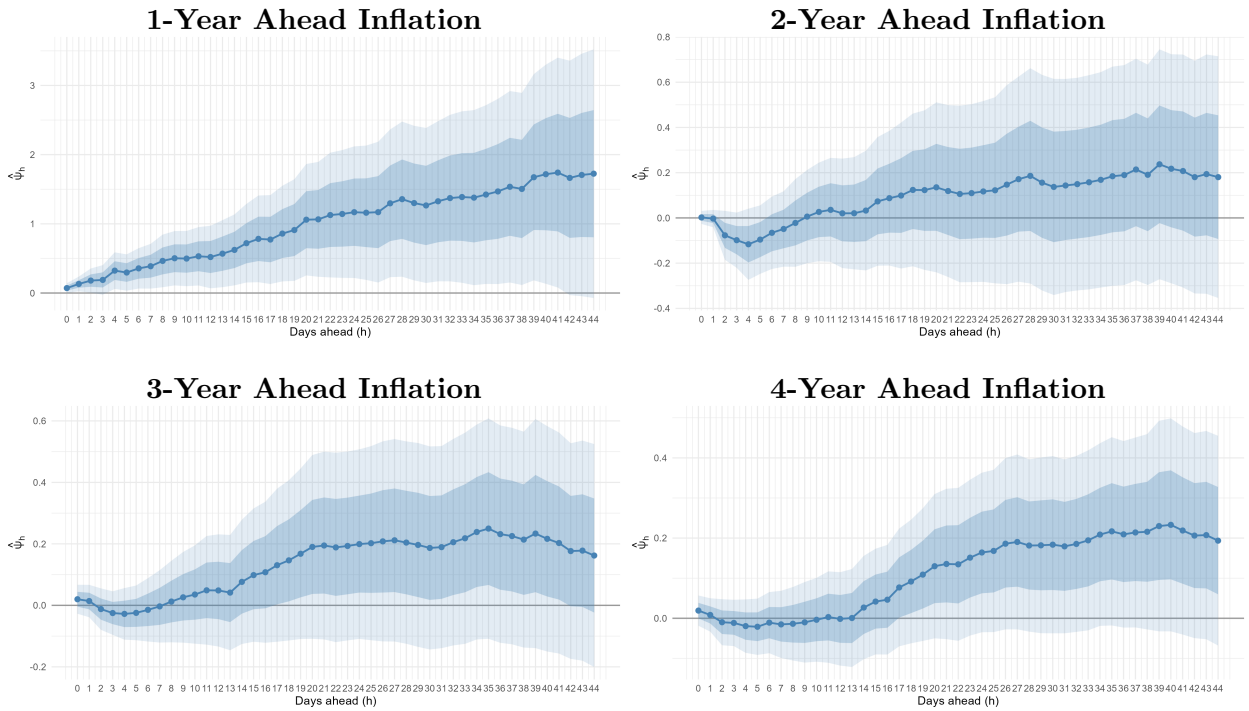


Figure 3. Estimated response of inflation expectations to a salient shock, by forecast horizon and trading day h. Top-left: 1-year ahead; top-right: 2-year ahead; bottom-left: 3-year ahead; bottom-right: 4-year ahead. Shaded areas report 68% and 90% bootstrap confidence bands.

This concentration is unsurprising: a major release is exactly the kind of public news on which institutions revise their forecasts and on which the financial press is most likely to report those revisions. But it is a concern for identification: on those days z_t moves with both the salient shock *and* an unusually large public macro innovation, and the heteroskedasticity in z_t across regimes may then capture a mixture of the two channels rather than the salient channel alone.

The objective of this subsection is to show that the heteroskedasticity estimator recovers the response to salient shocks $\eta_{t-1}^{(k_t)}$ rather than to public macro shocks ε_t. I implement three complementary corrections, and show that all three deliver an IRF essentially similar to the baseline of Figure 2, both in shape and in magnitude.

Controlling for individual release surprises. Because the Focus survey also reports institution-level forecasts of *monthly* inflation, I can construct, for each institution i and each release date, a daily time-series with individual release surprise

$$\text{Surprise}^{(i)}_{\text{IPCA},t} \equiv \mathbb{1}\{\text{ReleaseDate}^m = t\} \left(\pi_t^m - E^{(i)}_{t-1}[\pi_t^m]\right), \tag{31}$$

$$\text{Surprise}^{(i)}_{\text{IPCA-15},t} \equiv \mathbb{1}\{\text{ReleaseDate}^{15,m} = t\} \left(\pi_t^{15,m} - E^{(i)}_{t-1}[\pi_t^{15,m}]\right), \tag{32}$$

where m is the target month of the inflation release, and $ReleaseDate^m$ is the day in which the m month inflation was reported. The daily series is possibly non-zero only in days with IPCA or IPCA-15 monthly releases, which happens once a month in the beggining and middle of each month respectively. These are the genuine release surprises relative to each institution's own pre-release forecast and are by construction orthogonal to the information i used to form its prior forecast. I add both surprises (zero on non-release days) as additional regressors to the baseline LP (25),

$$E_t^{(i)}[\pi_{t+h}] = \alpha_h^{(i)} + \phi_h^m + \beta_{1,h} E_{t-1}^{(i)} + \beta_{2,h} \bar{E}_{t-1} + \beta_{3,h} \,\text{Surprise}_{\text{IPCA},t}^{(i)} + \beta_{4,h} \,\text{Surprise}_{\text{IPCA-15},t}^{(i)} + u_{t+h}^{(i)}, \quad (33)$$

and re-run the heteroskedasticity-based estimator on the new residuals. This regression leaves all R_1 days in the sample and remove the component of $u_{t+h}^{(i)}$ that is mechanically related to the realized release surprise. The resulting IRF is reported in Figure 4: it is virtually indistinguishable from the baseline, with the long-horizon point estimate slightly below 1.8 and the same monotone, hump-shaped profile.

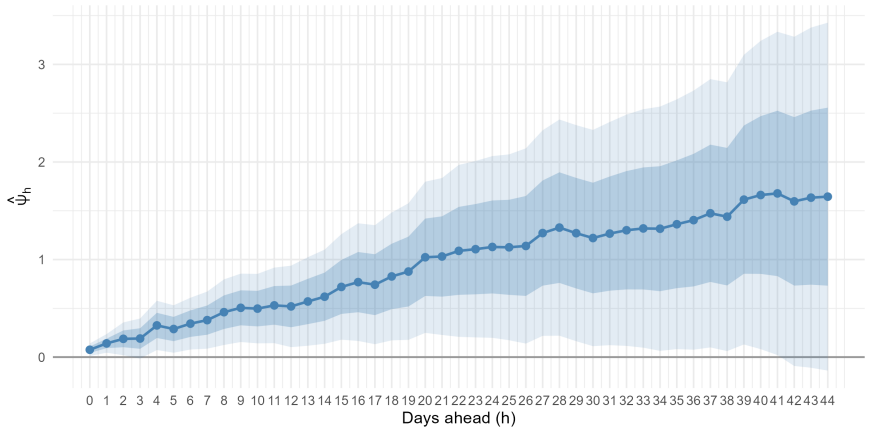


Figure 4. Response to a salient shock, controlling for individual IPCA and IPCA-15 release surprises in the baseline LP.

Orthogonalizing z_t from release surprises. Another test operates on the external variable rather than on the residual. I regress z_t on the cross-sectional average IPCA and IPCA-15 surprises,

$$z_t \;=\; \lambda_0 + \lambda_1 \,\overline{\text{Surprise}}_{\text{IPCA},t} + \lambda_2 \,\overline{\text{Surprise}}_{\text{IPCA-15},t} + z_t^{\perp}, \quad (34)$$

and use the residual $z_t^{\perp}$ as the external variable in the heteroskedasticity estimator. By construction $z_t^{\perp}$ is purged of the variation in the proxy attributable to the average realized release surprise; what remains is the variation due to all other channels, of which the salient channel on R_1 days is the relevant one. The between-regime heteroskedasticity of $z_t^{\perp}$ statistically significant, with a variance ratio $\sigma^2_{R_1}/\sigma^2_{R_2} = 1.44$ and a Brown–Forsythe, however weaker with F-statistic of $7,83$ ($p < 0.01$). Figure 5 reports the resulting IRF; once again, it is essentially identical to the baseline, with a

long-horizon plateau around 1.7 and the same weeks-long build-up.

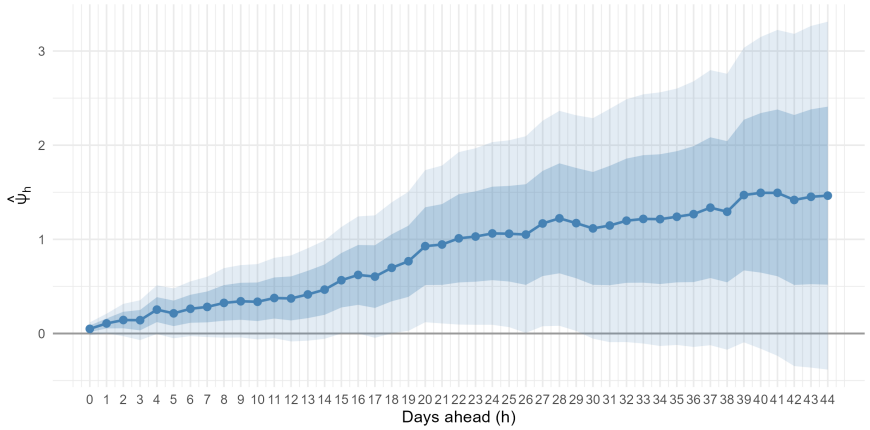


Figure 5. Response to a salient shock, using the orthogonalized external variable $z_t^{\perp}$.

Removing release days from R_1. Finally I drop from R_1 all days that coincide with an IPCA or IPCA-15 release. This shrinks R_1 from 258 to 151 days, but it removes by construction the contamination of (A2) from the most obvious source of public macro news. After this restriction the proxy z_t still exhibits a variance ratio of $\sigma^2_{R_1}/\sigma^2_{R_2} = 1.46$ — almost identical to the baseline — though the Brown–Forsythe test loses statistical significance ($F = 0.84$, $p = 0.36$), unsurprisingly given the roughly 40% loss of R_1 observations.The drop in test power reflects the loss of R_1 observationsThe estimator can still be applied to this restricted sample as a robustness check. Figure 6 reports the resulting IRF, which shows point estimates close to 5 and results not statistically significant even for 68% bands. Although not robust to this specification, such results could be pushed by the null power of the instrument.

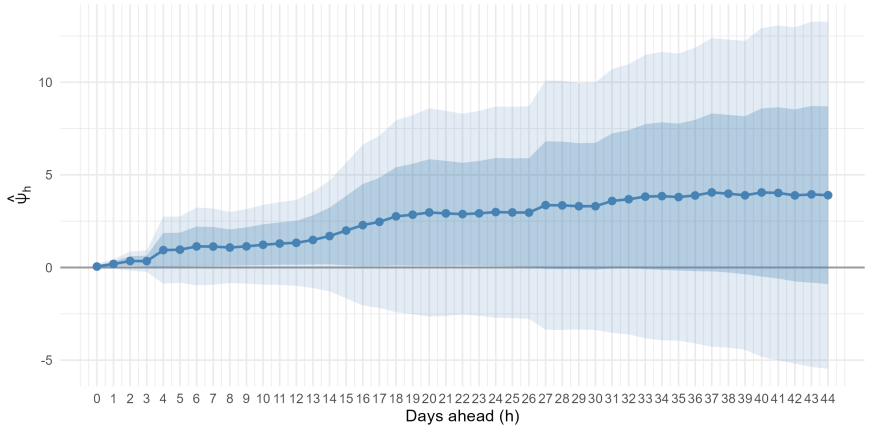


Figure 6. Response to a salient shock, removing IPCA and IPCA-15 release days from R_1.

What the three tests jointly show. Taken together, the three tests lead to the same conclusion. The baseline IRF and the IRFs from corrections (1)–(3) overlay each other in shape and magnitude: a near-zero on-impact response, a gradual build-up over the first six weeks of trading days, a crossing of unity at roughly three weeks, an apex of approximately two at seven weeks, and a long-horizon plateau. If the heteroskedasticity estimator were capturing a mixture of salient and public-macro channels, removing IPCA release-day variation from one side of the comparison or the other — be it the residual, the external variable, or the regime definition itself — should have moved the estimated response materially. It does not. I read this as strong evidence that the heteroskedasticity-based estimator is recovering the salient-channel response of (27) rather than the response to public macro shocks.

5 Conclusion

This paper started from a simple observation: professional forecasters do not form their views in isolation. Some institutions are prominent enough that the financial press reports their inflation forecasts, and when it does, the disclosed view carries the disclosing institution's idiosyncratic information — not just commonly available data. I formalized this as a *salient signal* that arrives lumpily through the media, embedded it in an otherwise standard model of strategic expectation formation, and showed that the channel turns one institution's private noise — a *salient shock* — into a disturbance to the entire cross-section of expectations, a granular mechanism in the spirit of Gabaix (2011). Taking the model to a confidential, institution-level panel of Brazilian inflation forecasts and identifying disclosure days with a large-language-model classifier, I used the heteroskedasticity of break-even inflation across disclosure and non-disclosure days to estimate the response of forecasters to a salient shock. The estimated impulse response is near zero on impact, consistent with the lumpy revision behavior of Baley and Turén (2025), builds gradually over the following weeks, and reaches a one-for-one response within roughly three weeks — direct evidence that forecasters do learn from their salient peers, and that this learning is slow rather than immediate.

These results come with clear limitations. The heteroskedasticity-based estimator recovers the response only up to the unknown loading of break-even inflation on the salient shock, so the point estimates are informative about the sign and shape of the response but not its exact level; the apparent overshooting at longer horizons should be read with this caveat in mind. Identification also relies on disclosure days not being systematically more volatile in macro news than other days. Two of my three robustness corrections leave the response essentially unchanged, but the most aggressive — dropping inflation-release days from the disclosure sample — loses statistical power, so I cannot fully rule out that release-day volatility contributes to the estimate. Finally, the salience weights are taken as exogenous and proxied by news coverage; endogenizing which institutions become newsworthy is left open.

The economic implication I want to emphasize is for monetary policy. The salient channel is a double-edged sword: it aggregates dispersed information, but it also amplifies one institution's idiosyncratic noise into the time-series volatility of aggregate expectations. Since aggregate inflation expectations are a central input to monetary policy — both as an object the central bank tries to anchor and as a signal it reads off market and survey data — a channel that raises their volatility is directly policy-relevant. If salient forecasters can move the consensus, then a central bank that reads aggregate expectations at face value may mistake a salient shock for a genuine shift in the inflation outlook, and respond to noise. This suggests that the disclosed views of prominent institutions deserve attention not only as forecasts but as a propagation channel that policy has to filter.

These implications can only be made precise within a general-equilibrium framework, which is the natural next step. The reduced-form impulse response estimated here can serve as a moment condition to recover the model's structural parameters via simulated method of moments, as in Gemmi and Valchev (2026), once the forecast-updating process is enriched with a simple friction allowing forecasters to revise their reported views only intermittently. With structural parameters in hand, one could embed the salient channel in a New Keynesian environment and ask three questions: how the strength of the channel reshapes the determinacy region of a Taylor rule, and thus how hawkish the central bank must be to keep expectations anchored; how much a salient shock moves output relative to a counterfactual economy without the channel — a granular monetary non-neutrality; and how a central bank that observes only aggregate expectations, unable to separate a salient shock from a fundamental one, should optimally filter its signal. These questions are left for my future work.

References

K. Adam, P. Kuang, and S. Xie. Overconfidence in private information explains biases in professional forecasts. May 2025. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=. Working paper.

I. Baley and J. Turén. Lumpy forecasts. December 23 2025. Unpublished manuscript.

C. Binder, J. R. Campbell, and J. M. Ryngaert. Consumer inflation expectations: Daily dynamics. *Journal of Monetary Economics*, 145:103613, May 2024. doi: 10.1016/j.jmoneco.2024.103613. URL https://doi.org/10.1016/j.jmoneco.2024.103613.

C. Binder, P. Frank, and J. M. Ryngaert. The causal effect of news on inflation expectations. Technical Report 34088, National Bureau of Economic Research, Aug. 2025. URL https://www.nber.org/papers/w34088.

P. Bordalo, N. Gennaioli, Y. Ma, and A. Shleifer. Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782, 2020. doi: 10.1257/aer.20181219.

O. Coibion and Y. Gorodnichenko. What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159, Feb. 2012. doi: 10.1086/665662. URL https://www.jstor.org/stable/10.1086/665662.

O. Coibion and Y. Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678, 2015. doi: 10.1257/aer.20110306.

G. W. Evans and S. Honkapohja. Adaptive forecasts, hysteresis, and endogenous fluctuations. *Economic Review*, (1):3–13, 1993.

X. Gabaix. The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772, May 2011. doi: 10.3982/ECTA8769.

L. Gemmi and R. Valchev. Biased surveys. *Journal of Monetary Economics*, 157:103868, 2026. ISSN 0304-3932. doi: 10.1016/j.jmoneco.2025.103868.

K. He, F. Sandomirskiy, and O. Tamuz. Private private information. *arXiv preprint arXiv:2112.14356*, Apr. 2025. URL https://arxiv.org/abs/2112.14356.

Z. Huo, M. Li, J. Wu, and H. Xiong. Micro shocks and macro fluctuations in the information network. Working paper, June 2025. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=.

Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, March 2005.

D. R. Känzig. The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review*, 111(4):1092–1125, 2021.

D. Laster, P. Bennett, and I. S. Geoum. Rational bias in macroeconomic forecasts. *The Quarterly Journal of Economics*, 114(1):293–318, Feb. 1999. doi: 10.1162/003355399555918.

A. Marcet and J. P. Nicolini. Recurrent hyperinflations and learning. *American Economic Review*, 93(5):1476–1498, Dec. 2003. doi: 10.1257/000282803322655400.

E. Nakamura and J. Steinsson. Identification in macroeconomics. *Journal of Economic Perspectives*, 32(3):59–86, Summer 2018.

M. Ottaviani and P. N. Sørensen. The strategy of professional forecasting. *Journal of Financial Economics*, 81(2):441–466, Aug. 2006. doi: 10.1016/j.jfineco.2005.08.002.

R. Rigobon. Identification through heteroskedasticity. *Review of Economics and Statistics*, 85(4): 777–792, 2003.

R. Rigobon and B. Sack. The impact of monetary policy on asset prices. *Journal of Monetary Economics*, 51(8):1553–1574, 2004.

E. Nakamura and J. Steinsson. High-Frequency Identification of Monetary Non-Neutrality: The Information Effect. *The Quarterly Journal of Economics*, 133(3):1283–1330, August 2018.

J. York. Do household inflation expectations respond to macroeconomic data releases? Job Market Paper, Nov. 2023. URL `https://jryork.com`.

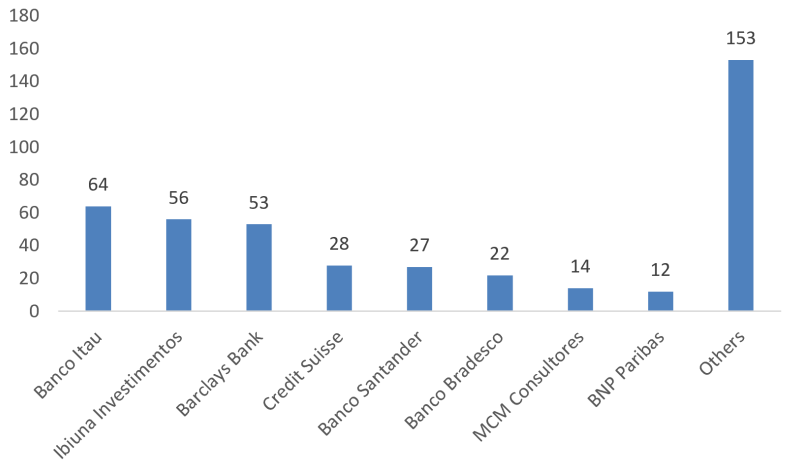


Figure 7. Number of news on inflation by institution (identified by LLM)

Appendix

A Riccati equation for the prior variance

The prior variance Σ evolves between periods according to

$$\Sigma \;=\; \rho^2\,\Sigma^{+} \;+\; \sigma_u^2, \tag{35}$$

where Σ^{+} is the posterior variance in the previous period. In a given regime r, the posterior variance is $\Sigma^{+,r} = \Sigma/(1+\Sigma\Pi^r) = \Sigma(1-G^{RE,r})$. Because each period is drawn i.i.d. between O and N with probabilities q and $1-q$, the steady-state prior variance must satisfy

$$\Sigma \;=\; \rho^2\,\Sigma\left[\frac{q}{1+\Sigma(\tau+\nu+\kappa)} \;+\; \frac{1-q}{1+\Sigma(\tau+\nu)}\right] \;+\; \sigma_u^2 \tag{36}$$

This is the *regime-averaged Riccati equation*. It pins down the steady-state prior variance as a function of the structural parameters τ, ν, ρ, σ_u^2, q, and the endogenous precision κ of the salient signal. Two boundary cases are illuminating. When $q = 1$, the salient signal arrives every period and (36) collapses to the baseline Riccati $\Sigma = \rho^2\Sigma/(1+\Sigma(\tau+\nu+\kappa)) + \sigma_u^2$ extended to a third signal. When $q = 0$, the salient channel is dormant and we recover the two-signal Riccati of **?**. In the empirically relevant intermediate region $q \in (0,1)$, Σ is a genuine convex combination, weighted by the *posterior variances* in each regime rather than by the precisions.

B LLM classification

B.1 Prompt for R_1 classification

1) You are an economist Reading News Headlines. Analyze the News below and decide whether the content from the News is about a financial institution revising their inflation expectations (or forecast) for Brazil. News can either be in english or portuguese

2) A news of expectation revision consists on a financial institution changing their forecast/expectation about the inflation rate (the same as IPCA, IGP-M ou INPC). The change in expectations can be either increases or decreases. Below I presente explicit examples:

- Bradesco revisa estimativas de IPCA, PIB e Selic : Here the institution Bradesco is revising their "estimations" which can be interpreted as expectations to IPCA which is the standard inflation index used in Brasil.

- Barclays eleva estimativa para IPCA 2021 de 5,7

- Credit Suisse eleva projecão IPCA 2021 de 6

- Itaú eleva projeção de IPCA em 2020 de 3,5

- Santander Sees Notable Impact on June CPI From Petrobras, Aneel: Also consider these cases in which one event explicitly triggering a forecast, in this case written as "sees notable impact".

- Bradesco BBI Cuts Brazil 2016 CPI Est. to 7

- Itau Cuts Estimates for Brazil GDP, Inflation and Trade Balance

Some advices, only consider the News that are very explicit about one specific private institution, as well as being directly related to revisions on inflation expectation. Bellow are some examples which, although ticky, you should NOT consider as a private institution revising expectations:

- Brazilian Stocks Gain on Interest Rate Signal, Slower Inflation: not considered because it is not highlighting na especific institution revising their expectation, and it does not indicate that the slower inflation is related to expectation or forecasts.

- Brazil Central Bank Signals End of Rate Increases: do not consider when the subject of the News is the Central Bank, and also this is talking about interest rates.

- IGP-M sobe 0,14

- FGV: IGP-M desacelera para 0,14

- Brazil Economists Forecast 2010 CPI at 5.18

3) Besides that, when you identify that the news is about expectations on inflation, I want you to identify which institution is making the revision, here is a list of possible institutions.

Banco ABC Brasil; BNY Mellon ARX Investimentos; Bahia AM; Banrisul; Banco do Brasil; Haitong Banco de Investimento; Banco BNP Paribas; Banco Bradesco; BB DTVM; Banco Citibank; Deutsche Bank - Banco Alemao S.A.; Banco Fibra; GAP AM; Verde Asset Management; Ventor Investimentos; Banco Itau; LCA Consultores; MB Associados; MCM Consultores; Bank of America Merrill Lynch Banco Multiplo; Banco MUFG; Opportunity Asset Management; BTG Pactual AM; Banco Pine; Previ; Rosenberg e Associados; Banco Ribeirao Preto; Banco Safra; Banco Santander;

Telefonica / Vivo; Tendências Consultoria Integrada; Votorantim Asset; Banco Mizuho; Modal Asset Management; Quest Investimentos Ltda.; Claritas Adm Rec; Icatu Vanguarda Administracao de Recursos; Petros Fundacao de Seguridade Social; Maua Investimentos; Banco Sicredi; Ativa Investimentos; Funcef - Fund. Economiarios Federais; Pezco Economics; Link S/A CCTVM; Brasilprev; Bradesco - Asset Management; Gradual Investimentos; Kondor Adm Gestora; Paineiras Investimentos; Rio Bravo; BNP Paribas AM; JGP Gestao de Recursos; M. Safra Co; CM Capital Markets; Banco Fator; Sul America Investimentos; Serasa Experian; Porto Seguro Investimentos; Vinci Gestora; Kapitalo Investimentos; Votorantim Asset; Franklin Templeton Investimentos Brasil; SPX Capital; Planner Corretora de Valores; Brasil Plural Capital Gestao de Recursos; BW Gestao; Medley Global Advisors; Banco Morgan Stanley; Santander AM; Ibiuna Investimentos; Caixa Asset; Rabobank; Quantitas Asset Management; Nobel AM; XP Gestao; Canvas Capital; FGV; Absolute Gestao; Garde AM; Somma Investimentos; Vintage Investimentos; Itaim AM; Parcitas Gestao de Investimentos; Parallaxis Consultoria Economica; Bahia AM; 4E Consultoria; Reag Investimentos; Empiricus Consultoria; Daycoval AM; Ethica Asset Management; Canepa AM; Mongeral Aegon Inv; Rio Gestão de Recursos; BGC Liquidez; Truxt Investimentos; Guide Investimentos; Mogno Capital Investimento Ltda; Barclays Bank; MZK Investimentos;

There can be more institutions besides these quoted in here, also you can see some slight variations of the names.

4) Here is the news, analyse:(include news headline in here)