Misperceived Propagation in Macroeconomic Expectations*

Zhesheng Qiu†

First version: July 03, 2023.
This version: September 14, 2026.

Abstract

Survey forecasts can misperceive the present or attach the wrong future path to it. This paper identifies the second with the *propagation wedge*, the gap between actual and perceived propagation, which is zero under correct propagation whatever forecasters observe or how much they trust it. U.S. survey forecasts under-propagate changes in unemployment and the short-term interest rate but over-propagate inflation. In the forecast errors that current outcomes predict, the wedges are as large as the information-friction component. Projection bias, a pull toward the current level, produces both signs. Federal Reserve staff forecasts carry the same misperceptions.

Keywords: Economic forecasts; misperceived propagation; monetary policy.
JEL Classification: D84, D91, E17, E27, E37, E71.

*I thank George-Marios Angeletos, Benjamin Born, Olivier Coibion, Wei Cui, Áureo de Paula, Ding Dong, Joel Flynn, Xavier Gabaix, Chenyu Hou, Zhen Huo, Dun Jia, Pei Kuang, Laura Liu, Yulei Luo, Yueran Ma, Pooya Molavi, Emi Nakamura, Kristoffer Nimark, Guangyu Pei, José-Víctor Ríos-Rull, Dmitriy Sergeyev, Alireza Tahbaz-Salehi, Takayuki Tsuruga, Harald Uhlig, Christian Wolf, Donghoon Yoo, and Eric Young for helpful discussion at different stages of this paper. I also thank participants at the WESEAMS meeting, AEA meeting, the HKU–UCL–ESRC Summer Workshop, Hong Kong Junior macro group meeting at CUHK, HKUST, Osaka University, and Shandong University for their comments. Qing Wang and Yu Zhang provided excellent research assistance. Financial support from the CityU Start-up Grant 7200696 is gratefully acknowledged. All remaining errors are my own.

†Department of Economics, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong. Email: zheshqiu@ust.hk.

Forecasters can misperceive either the current state of the economy or its future propagation. Much of the macroeconomic effect of expectations operates through this latter margin, because households, firms, and central banks all make decisions based on their forecasts of the future. Better information about the present corrects the first margin and leaves the second untouched. A growing class of models locates the departure from rational expectations in this margin. For example, forecasters under-propagate their nowcast under cognitive discounting (Gabaix 2020) and over-propagate it under over-extrapolation (Angeletos, Huo and Sastry 2020). However, when the true data-generating process is unknown, standard empirical diagnostics applied to survey data do not disentangle the propagation margin from other margins. This paper proposes a new diagnostic statistic that does—the propagation wedge.

The wedge is the gap between the coefficients of two mirror-image instrumental-variables regressions. Actual propagation is captured by regressing the future outcome on the current realization, instrumenting the realization with the nowcast; perceived propagation is captured by regressing the forecast of the future outcome on the nowcast, instrumenting the nowcast with the current realization. The wedge takes a simple form in two special cases: in an AR(1) environment, it reduces to the gap between actual and perceived persistence, and under full information it equals the coefficient from regressing forecast errors on current realizations.

In a general environment with multidimensional dynamics and imperfect information, the wedge has three advantages. First, a nonzero wedge rejects rational expectations without the econometrician specifying what forecasters know. Information rigidity, rational inattention, and imprecise memory are all rational expectations on limited information, so they sit inside the zero-wedge null.[1] Second, under linear forecasting rules, a nonzero wedge further rules out a misperception of signal precision alone (in the sense of observational equivalence). Therefore, the diagnostic Kalman filter, ambiguity aversion about the mean, and overconfidence in private information[2] also sit inside the zero-wedge null. Third, the sign of the wedge separates under-propagation from over-propagation,

[1] See Mankiw and Reis (2002); Maćkowiak and Wiederholt (2009); da Silveira, Sung and Woodford (2024).
[2] See Bordalo et al. (2020); Huo, Pedroni and Pei (2024); Adam, Kuang and Xie (2025).

while its size can be set against information frictions in accounting for the predictable part of forecast errors. The first two advantages complement the diagnostic regressions of Coibion and Gorodnichenko (2015) and Bordalo et al. (2020), whose nulls lie inside the wedge's, while the last one characterizes and quantifies misperceived propagation.

In the U.S. Survey of Professional Forecasters (SPF), the present is carried forward incorrectly, and the error has a sign. From 1984 to 2018, consensus forecasts, the mean across forecasters, carried forward about half of the actual propagation of a change in the unemployment rate and less than a quarter of that of a change in the 3-month Treasury-bill rate, so movements in the two rates die out too fast in the forecasts. They carried forward about one and a half times the actual propagation of CPI inflation, so inflation is carried too far. Averaged over one to four quarters ahead, the wedge is $+0.24$ for the unemployment rate, $+0.21$ for the Treasury-bill rate and -0.13 for inflation, with standard errors of 0.10, 0.06 and 0.06 (Figure 1). The sign holds at every horizon, in every fifteen-year window of the sample, in the individual responses as in the consensus, and in the Livingston Survey (Figure 2). The Federal Reserve Board staff's Greenbook projections carry the same signs, at sizes close to the survey's for the unemployment rate and inflation. The staff is a single institutional forecaster, so no survey average produces the pattern. Its projections are what the Federal Open Market Committee weighs when it sets the policy rate, so the misperception enters the policy process at its input.

The misperception is as large as the information friction in the same forecast errors. The consensus forecast error has a part that the current outcome predicts, and that part splits exactly in two: the wedge, and an *information-friction component* that vanishes when the nowcast equals the realized present. For the two rates the two parts are of the same size; for inflation the friction component is indistinguishable from zero, and the wedge is the whole of a small predictable error (Figure 3). That predictable part is a small share of the errors' variance, so the margin governs how forecasts respond to the present rather than how large their errors are. Information rigidity cannot produce these numbers. In an economy calibrated to the three series, in which every forecaster is rational and perceives propagation correctly but the consensus under-reacts to news, the wedge is exactly zero. The two regressions that ask whether forecast errors are predictable from the present are

not: they reach 30 to 45 percent of their SPF values for the two rates, and for inflation they take the sign opposite to the data.

Which account of expectation formation produces the pattern? Because the three series carry both signs, any bias that moves every forecast the same way fails. Cognitive discounting under-propagates every series whose actual propagation is positive, and so predicts a positive wedge for inflation, whose measured wedge is negative. Over-extrapolation, natural expectations (Fuster, Laibson and Mendel 2010), diagnostic expectations that leave the nowcast undistorted (Bordalo, Gennaioli and Shleifer 2018), simple models (Molavi 2025) and constant-gain learning that has settled (Evans and Honkapohja 2001) each predict a sign, or a wedge at or near zero, that at least one of the three series contradicts (Table 2). What survives is a *level anchor*, a pull toward the present, whose simplest form is projection bias, the tendency to expect current conditions to persist (Loewenstein, O'Donoghue and Rabin 2003). A forecast pulled toward the current level carries every level too far and, because the pull cancels when the forecast is differenced, damps every change: the sign follows the form of the series, level or change, rather than the variable. A belief that part of the current level is permanent produces the same pattern, so the data speak to the anchor, not to projection bias in particular.

The anchor is testable beyond the three signs, because it is one parameter per series, the weight on the present, and the wedge it implies is that weight times a factor set by actual propagation alone. Each headline series can be measured in both forms, and the weight fitted on the form reported above predicts the sign of the wedge in the other form with no free parameter (Figure 4). The prediction holds in all three series, and in all nine series for which both forms can be built. Across the twenty-nine cells the survey supports, each series in each form it allows, the weight pooled over the three headline series gives the sign of the wedge in 22 of the 26 cells it was not fitted on. The anchor also reaches the benchmark regression of the survey literature, of the forecast error on the forecast revision: under the maintained model the anchor contributes to its slope, and for the unemployment rate the slope net of the anchor is not distinguishable from zero at any horizon, though with wide intervals (Figure 5).

The same separation of weight from factor holds over time. On rolling fifteen-year windows, actual propagation moves, the wedge moves with it, and the weight moves far less (Figure 7). When the 2020 quarters enter the unemployment windows, actual propagation turns negative and the wedge changes sign with it; the implied weight does not. As the actual propagation of inflation falls from 0.4 to about zero across the sample, perceived propagation falls with it, and the weight is unchanged. Run back as far as each survey question reaches, the weight drifts the same way in all three series, downward over forty years, and most for the unemployment rate.

Related literature. The paper belongs to the work that reads the formation of expectations off survey forecasts. Early tests asked whether forecasts are unbiased and efficient (Mincer and Zarnowitz 1969; Keane and Runkle 1990). Lovell (1986) showed why their rejections are hard to read: a test of rationality is also a test of the information set the econometrician assigns to forecasters, the joint hypothesis the wedge's null avoids. Coibion and Gorodnichenko (2015) derived from models of sticky and noisy information a regression of consensus forecast errors on forecast revisions, whose coefficient is zero under full-information rational expectations and positive under information rigidity. Bordalo et al. (2020) run the regression forecaster by forecaster and read the overreaction they find as diagnostic expectations. The later readings of these regressions (Kohlhas and Walther 2021; Sung 2026; Huo, Pedroni and Pei 2024; Adam, Kuang and Xie 2025) are the models listed above, and every one of them sits inside the wedge's null (Table 1). The wedge's nearest neighbors read perceived persistence off survey forecasts, Reis (2020) from long- and short-horizon forecasts and Angeletos, Huo and Sastry (2020) within an economy that follows an AR(1); the wedge requires no assumption on the law of motion.

The paper also relates to the models of expectation formation built for macroeconomic use, which the wedge sorts. The behavioral inattention surveyed by Gabaix (2019) differs from rational inattention on exactly the wedge's line: the forecaster acts on a misspecified model of the economy rather than on a conditional expectation, and the misspecification opens a wedge once it reaches the dynamics, as in cognitive discounting. Other models that distort how the present is carried forward include, besides the accounts the three se-

ries reject, horizon-specific diagnosticity (Bianchi, Ilut and Saijo 2026) and learning while beliefs are still moving (Eusepi and Preston 2011; Farmer, Nakamura and Steinsson 2024). Projection bias was proposed for predictions of one's own future tastes (Loewenstein, O'Donoghue and Rabin 2003); here a pull of that form is fitted to professional forecasts of the macroeconomy, with a weight read from how the forecasts comove with the present rather than from their errors. The null also covers forecasts that condition on news about the future (Goldstein and Gorodnichenko 2023), because a conditional expectation on any information set produces a zero wedge.

The paper proceeds as follows. Section 1 constructs the wedge, proves its null and places the leading models of information and precision inside it. Section 2 measures the wedge on the three headline series and across forecasters, sets it against the information-friction component, and compares it with the alternative diagnostics. Section 3 derives the sign each leading account implies, fits one weight per series, and carries that weight into the series' other form, the error–revision regression, the other series the survey asks and every sample window. Section 4 concludes.

1 The propagation wedge

This section develops the *propagation wedge*, a diagnostic of the gap between how a macroeconomic object actually propagates and how forecasters perceive it. The wedge draws on neither the forecasters' information sets nor a model of the economy's dynamics.

1.1 Constructing the wedge

Throughout, let z be the variable of interest, and $\mathbb{F}_t$ the time-t *consensus* forecast operator—the cross-sectional mean of the *individual* forecasts $\mathbb{F}_{i,t}$, taken over the same forecasters at the nowcast and at every horizon. Thus, $\mathbb{F}_t z_t$ and $\mathbb{F}_t z_{t+h}$ are the consensus nowcast and h-step forecast. Under information frictions, $\mathbb{F}_t z_t$ need not equal z_t. As a regularity condition, all second moments of z and of the forecasts are assumed finite.

A natural diagnostic of whether forecasters carry the present into the future correctly

would compare two persistence regressions: the slope of z_{t+h} on z_t against that of $\mathbb{F}_t z_{t+h}$ on $\mathbb{F}_t z_t$. This comparison may misread information rigidity as propagation misperception when the true model is not AR(1).

Example 1 (Naive persistence comparison in an AR(2) model)**.** Let $z_{t+1} = \phi_1 z_t + \phi_2 z_{t-1} + \epsilon_{t+1}$ with i.i.d. ϵ_{t+1}, and let forecasters hold rational expectations on the lagged information set z^{t-1}, so that $(\mathbb{F}_t z_{t-1}, \mathbb{F}_t z_t) = (z_{t-1}, z_t - \epsilon_t)$, and

$$\frac{\operatorname{cov}(z_{t+1},\, z_t)}{\operatorname{var}(z_t)} = \phi_1 + \phi_2 \frac{\operatorname{cov}(z_{t-1},\, z_t)}{\operatorname{var}(z_t)}, \quad \frac{\operatorname{cov}(\mathbb{F}_t z_{t+1},\, \mathbb{F}_t z_t)}{\operatorname{var}(\mathbb{F}_t z_t)} = \phi_1 + \phi_2 \frac{\operatorname{cov}(z_{t-1},\, z_t)}{\operatorname{var}(\mathbb{F}_t z_t)}. \quad (1)$$

In Example 1, forecasters perceive the AR(2) process correctly, yet the two slopes differ: the gap is proportional to ϕ_2 and reflects the information lag between the outcome z_t and its nowcast $\mathbb{F}_t z_t$. The repair is to construct the propagation wedge as in Definition 1. The information friction enters the two ratios alike and hence cannot by itself open a gap, regardless of the data-generating process (proved in Section 1.2).

Definition 1 (Actual and perceived propagation, and the wedge)**.** At horizon h, *actual* and *perceived* h-step propagation are

$$a_h \equiv \frac{\operatorname{cov}(z_{t+h},\, \mathbb{F}_t z_t)}{\operatorname{cov}(z_t,\, \mathbb{F}_t z_t)}, \qquad \widetilde{a}_h \equiv \frac{\operatorname{cov}(\mathbb{F}_t z_{t+h},\, z_t)}{\operatorname{cov}(\mathbb{F}_t z_t,\, z_t)}, \quad (2)$$

and the *propagation wedge* is the gap between them,

$$\tau_h \equiv a_h - \widetilde{a}_h. \quad (3)$$

Forecasters *under-propagate* z when $\tau_h > 0$ and *over-propagate* it when $\tau_h < 0$.[3]

Each ratio is an instrumental-variables regression slope: a_h from regressing z_{t+h} on z_t with the nowcast as instrument, and $\widetilde{a}_h$ from regressing $\mathbb{F}_t z_{t+h}$ on $\mathbb{F}_t z_t$ with the realized present as instrument. Each therefore reads how the present *propagates* into the future—a_h in the data and $\widetilde{a}_h$ in the forecasts.

[3]Propagation here is the predictive link from the present to future outcomes, not an impulse response.

Because covariance is symmetric in its arguments, the two slopes in equation (2) share the denominator $\operatorname{cov}(\mathbb{F}_t z_t,\, z_t)$, assumed nonzero throughout; and the wedge vanishes if and only if $\operatorname{cov}(z_{t+h},\, \mathbb{F}_t z_t) = \operatorname{cov}(\mathbb{F}_t z_{t+h},\, z_t)$. In the environment of Example 1, both of the covariance terms equal $\phi_1 \left[\operatorname{var}(z_t) - \operatorname{var}(\epsilon_t)\right] + \phi_2 \operatorname{cov}(z_{t-1},\, z_t)$ at $h = 1$, so the wedge stays silent exactly where the naive persistence comparison misfires. The following two examples show what a nonzero wedge captures in polar cases.

Example 2 (Univariate state)**.** Let $z_{t+h} = \rho^h z_t + \xi_{t+h}$ with ξ_{t+h} orthogonal to the present outcome z_t and to the nowcast $\mathbb{F}_t z_t$. When all forecasters iterate their nowcasts at a single perceived persistence $\widetilde{\rho}$, the forecast satisfies $\mathbb{F}_t z_{t+h} = \widetilde{\rho}^h \mathbb{F}_t z_t$, and

$$a_h = \rho^h, \qquad \widetilde{a}_h = \widetilde{\rho}^h, \qquad \tau_h = \rho^h - \widetilde{\rho}^h. \tag{4}$$

In Example 2, the wedge exactly measures the misperception of persistence, regardless of the information structure.[4] The headline settings of Coibion and Gorodnichenko (2015) and Bordalo et al. (2020) are versions of this example: information rigidity and diagnostic expectations distort the nowcast, but both carry the nowcast forward at the true ρ, so $\widetilde{\rho} = \rho$ and the wedge is silent for both. Perceiving ρ too low or too high is what Angeletos, Huo and Sastry (2020) call under- and over-extrapolation; Definition 1 names these directions under- and over-propagation, because in more general environments no single parameter of extrapolation can summarize a misperception that runs across the whole future path.

Example 3 (Full information)**.** Let z follow *any* process, and suppose forecasters observe the present perfectly, so $\mathbb{F}_t z_t = z_t$ and $\operatorname{cov}(\mathbb{F}_t z_t, z_t) = \operatorname{var}(z_t)$. Then

$$a_h = \frac{\operatorname{cov}(z_{t+h}, z_t)}{\operatorname{var}(z_t)}, \qquad \widetilde{a}_h = \frac{\operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t)}{\operatorname{var}(z_t)}, \qquad \tau_h = \frac{\operatorname{cov}(z_{t+h} - \mathbb{F}_t z_{t+h},\, z_t)}{\operatorname{var}(z_t)}. \tag{5}$$

In Example 3, the wedge τ_h is the slope of the forecast error $z_{t+h} - \mathbb{F}_t z_{t+h}$ on the current outcome z_t, the coefficient that Kohlhas and Walther (2021) call extrapolation.

These two polar cases give the wedge a transparent interpretation, but each obtains it by shutting down one source of complexity: multidimensional dynamics in the first, and

[4] Reis (2020) reads $\widetilde{\rho}$ off the long- and short-horizon survey forecasts in the same spirit.

imperfect information in the second. The next subsection asks what the wedge identifies in general environments when both are present.

1.2 What a nonzero wedge identifies

In general environments—with multidimensional dynamics and imperfect information—a nonzero wedge rejects rational expectations. Under covariance stationarity and linear forecasting rules, it further identifies what is misperceived: propagation.

Proposition 1 (The rational-expectations null)**.** *Suppose every forecaster i reports her rational forecasts, i.e., the conditional expectations on her information set $\mathcal{I}_{i,t}$,*

$$\mathbb{F}_{i,t} z_{t+h} = \mathbb{E}[z_{t+h} | \mathcal{I}_{i,t}], \qquad h = 0, 1, \ldots. \tag{6}$$

Then $\tau_h = 0$ for every $h \geq 1$.

Proof. For any square-integrable X and any square-integrable $\mathcal{I}_{i,t}$-measurable W, the law of iterated expectations gives $\operatorname{cov}(X, W) = \operatorname{cov}(\mathbb{E}[X|\mathcal{I}_{i,t}], W)$. Applying this identity to $X = z_{t+h}$ and then to $X = z_t$ yields

$$\operatorname{cov}\big(z_{t+h}, \mathbb{E}[z_t|\mathcal{I}_{i,t}]\big) = \operatorname{cov}\big(\mathbb{E}[z_{t+h}|\mathcal{I}_{i,t}], \mathbb{E}[z_t|\mathcal{I}_{i,t}]\big) = \operatorname{cov}\big(\mathbb{E}[z_{t+h}|\mathcal{I}_{i,t}], z_t\big),$$

in which the two outer terms are $\operatorname{cov}(z_{t+h}, \mathbb{F}_{i,t} z_t)$ and $\operatorname{cov}(\mathbb{F}_{i,t} z_{t+h}, z_t)$ under equation (6). Averaging them over i and dividing them by the shared denominator $\operatorname{cov}(\mathbb{F}_t z_t, z_t)$ gives $a_h = \widetilde{a}_h$ and hence $\tau_h = 0$. □

Proposition 1 holds for any law of motion and any information structure, with no linearity, distributional, or stationarity requirement.[5] Sticky information, noisy information, rational inattention, and imprecise memory models all represent each forecast as a conditional expectation on some information set, so Proposition 1 places them inside the null

[5]The proof uses only the law of iterated expectations. The covariance equality behind the null therefore holds date by date without stationarity, and for any information set, whether it looks backward or forward. In Goldstein and Gorodnichenko (2023), for example, forecasts condition on news about the future as well, and Proposition 1 applies unchanged.

at once.[6]

A nonzero wedge therefore rejects rational expectations in the sense of Muth (1961)—without requiring the econometrician to specify what the forecasters know. The rejection escapes the joint-hypothesis problem of Lovell (1986): a conventional rationality test rests on a correct guess of the forecasters' information set.[7] Further identifying what forecasters misperceive, rather than merely rejecting rational expectations, requires structure on their forecasting rules.

Assumption 1 (Linearity and stationarity)**.** For each forecaster i, there is a finite-dimensional signal vector $\mathbf{S}_{i,t} \in \mathbb{R}^{n_i}$ such that

$$\mathbb{F}_{i,t} z_{t+h} = \boldsymbol{\beta}_{i,h}^{\top} \mathbf{S}_{i,t}, \qquad h = 0, 1, \ldots, \tag{7}$$

with the loadings $\{\boldsymbol{\beta}_{i,h}\}$ non-random and constant over time. The joint process $(\mathbf{S}_{i,t}, z_t)$ is covariance stationary, with signal–outcome covariance $\mathbf{C}_{i,h} \equiv \operatorname{cov}(\mathbf{S}_{i,t}, z_{t+h})$ and positive-definite signal variance $\boldsymbol{\Sigma}_i \equiv \operatorname{var}(\mathbf{S}_{i,t})$.

The linearity in equation (7) only restricts the forecasting rule. It does not restrict what the forecaster looks at, how the loadings on signals differ across forecasters and horizons, or where the loadings come from.[8] What binds is that the loadings do not respond to any realized outcomes, the restriction that Section 1.4 relaxes for learning models. Stationarity makes the wedge a single number at each horizon; without it, the wedge is date-specific.

A forecaster who correctly perceives $\boldsymbol{\Sigma}_i$ and $\{\mathbf{C}_{i,h}\}$ loads $\boldsymbol{\Sigma}_i^{-1}\mathbf{C}_{i,h}$ and reports the best linear unbiased predictor—the *actual-BLUP*. If instead she perceives a symmetric positive-definite signal variance $\widetilde{\boldsymbol{\Sigma}}_i$ and signal–outcome covariances $\{\widetilde{\mathbf{C}}_{i,h}\}$, her loadings are $\widetilde{\boldsymbol{\Sigma}}_i^{-1}\widetilde{\mathbf{C}}_{i,h}$. A forecasting rule with such a representation is called a *perceived-BLUP*.

Proposition 2 (Completeness and identification)**.** *Let Assumption 1 hold.*

[6] Mankiw and Reis (2002), Woodford (2003), Maćkowiak and Wiederholt (2009), and da Silveira, Sung and Woodford (2024), respectively. In the last, recall is itself noisy, so the forecast is a conditional expectation on degraded signals rather than on a subset of what was once known.

[7] The information-set guess can fail quietly: when memory is imprecise, a forecast revision may also arise from degraded past signals and therefore need not proxy for newly received information (see Sung 2026).

[8] All series are demeaned, so the forecasting rules carry no intercept—a normalization, not a restriction, since constants enter no covariance. The collinear signals are dropped, so the forecast loadings are unique.

(i) *Every individual forecasting rule* $\{\boldsymbol{\beta}_{i,h}\}$ *admits a perceived-BLUP representation, including one with* $\widetilde{\boldsymbol{\Sigma}}_i = \boldsymbol{\Sigma}_i$.

(ii) *Fix a horizon* $h \geq 1$. *If, for every forecaster* i, *there exists a perceived-BLUP representation satisfying* $\widetilde{\mathbf{C}}_{i,0} = \mathbf{C}_{i,0}$ *and* $\widetilde{\mathbf{C}}_{i,h} = \mathbf{C}_{i,h}$, *then* $\tau_h = 0$.

Proof. For part (i), set $\widetilde{\boldsymbol{\Sigma}}_i = \boldsymbol{\Sigma}_i$ and $\widetilde{\mathbf{C}}_{i,h} = \boldsymbol{\Sigma}_i \boldsymbol{\beta}_{i,h}$, so that $\widetilde{\boldsymbol{\Sigma}}_i^{-1} \widetilde{\mathbf{C}}_{i,h} = \boldsymbol{\beta}_{i,h}$. For part (ii), the perceived loadings satisfy $\boldsymbol{\beta}_{i,\ell}^\top = \widetilde{\mathbf{C}}_{i,\ell}^\top \widetilde{\boldsymbol{\Sigma}}_i^{-1} = \mathbf{C}_{i,\ell}^\top \widetilde{\boldsymbol{\Sigma}}_i^{-1}$ for $\ell \in \{0, h\}$, so that

$$\operatorname{cov}(\mathbb{F}_{i,t} z_{t+h}, z_t) = \mathbf{C}_{i,h}^\top \widetilde{\boldsymbol{\Sigma}}_i^{-1} \mathbf{C}_{i,0}, \qquad \operatorname{cov}(z_{t+h}, \mathbb{F}_{i,t} z_t) = \mathbf{C}_{i,0}^\top \widetilde{\boldsymbol{\Sigma}}_i^{-1} \mathbf{C}_{i,h}. \tag{8}$$

The two right-hand sides are scalar transposes of each other, hence equal because $\widetilde{\boldsymbol{\Sigma}}_i^{-1}$ is symmetric. Averaging them across forecasters and dividing them by a common $\operatorname{cov}(\mathbb{F}_t z_t, z_t)$ gives $a_h = \widetilde{a}_h$, i.e., $\tau_h = 0$. □

Part (i) establishes that every linear forecasting rule minimizes the mean squared error under some perceived moments—an observational equivalence, not an assumption. The wedge therefore does not ask whether forecasters optimize; it asks whether the perceived moments can be the true ones. For the perceived variance the answer is always yes, so a misperceived variance is not testable either.

For the signal–outcome covariance, the answer depends on the rule. With the loadings fixed by the rule, each choice of $\widetilde{\boldsymbol{\Sigma}}_i$ implies a perceived covariance $\widetilde{\mathbf{C}}_{i,h} = \widetilde{\boldsymbol{\Sigma}}_i \boldsymbol{\beta}_{i,h}$. A single choice must therefore make the covariance with the present and that with the future both correct at once; some rules permit this. Part (ii) places them inside the null: for them, the perceived variance $\widetilde{\boldsymbol{\Sigma}}_i$ affects the two terms of equation (8) in the same way, and $\tau_h = 0$.

A nonzero wedge places at least one forecaster outside the null: no choice of her perceived variance makes both $\widetilde{\mathbf{C}}_{i,0}$ and $\widetilde{\mathbf{C}}_{i,h}$ correct. Since $\widetilde{\mathbf{C}}_{i,0}$ determines her nowcast and $\widetilde{\mathbf{C}}_{i,h}$ her h-step forecast, the pair determines how she carries the present forward. That is what the wedge finds wrong: *misperceived propagation*, the attribution that linearity buys.

Example 4 shows why the leading model of overreaction, the diagnostic Kalman filter of Bordalo et al. (2020), violates rational expectations without opening a wedge.

Example 4 (Diagnostic Kalman filter)**.** Let z be an AR(1) with persistence ρ, where $|\rho| < 1$, and let $x_{i,t}$ denote forecaster i's rational nowcast—the steady-state Kalman-filter estimate—with variance $V \equiv \text{var}(x_{i,t})$. She overweights the period-t news by a factor of $\theta > -1$ and carries the result forward by the actual persistence,

$$\mathbb{F}_{i,t} z_t = x_{i,t} + \theta\big(x_{i,t} - \rho\, x_{i,t-1}\big), \qquad \mathbb{F}_{i,t} z_{t+h} = \rho^h\, \mathbb{F}_{i,t} z_t. \tag{9}$$

The rule is linear in $x_{i,t}$ and $x_{i,t-1}$, both her conditional expectations, so $\mathbf{S}_{i,t} \equiv (x_{i,t},\, x_{i,t-1})^\top$ and $\boldsymbol{\beta}_{i,h} = \rho^h \cdot (1+\theta,\, -\theta\rho)^\top$. A perceived-BLUP keeping the signal–outcome covariance profile correct, $\widetilde{\mathbf{C}}_{i,h} = \mathbf{C}_{i,h}$, carries the perceived signal variance

$$\widetilde{\boldsymbol{\Sigma}}_i = V \begin{pmatrix} \frac{1+\theta\rho^2}{1+\theta} & \rho \\ \rho & 1 \end{pmatrix}, \tag{10}$$

a valid covariance matrix whose leading entry falls below V when $\theta > 0$ (Appendix A1).

In Example 4, the perceived signal–outcome covariance is correct, so Part (ii) of Proposition 2 delivers $\tau_h = 0$: the model sits inside the null despite violating rational expectations.[9] In the representation, her overreaction is a misreading of her nowcast's precision—the leading entry of equation (10) falls below V.

In sum, a nonzero wedge rejects rational expectations. Under Assumption 1, it further locates the misperception in the signal–outcome covariance profile, which corresponds to propagation. Section 1.3 reads leading models of expectations in the literature against the wedge, and compares it with two other diagnostic statistics already in wide use.

1.3 Zero wedge in leading models

A large class of leading expectations models has nonzero coefficients in the diagnostics of Coibion and Gorodnichenko (2015) (CG) or Bordalo et al. (2020) (BGMS) but still a zero

[9]The representation requires the misperception to be uniform across horizons. Otherwise, the misperception cannot be attributed to the variance of signals shared by all horizons, and has to be rationalized by misperceived horizon-specific signal–outcome covariance. The diagnostic expectations of Bordalo, Gennaioli and Shleifer (2018) leave the nowcast undistorted, while the smooth diagnostic expectations of Bianchi, Ilut and Saijo (2026) make diagnosticity horizon-specific; both departures can open a wedge.

propagation wedge (Table 1). Those zeros are Propositions 1 and 2 applied to each model: the table's blocks sort the models by what they distort—nothing, what forecasters observe (information), and how much they trust it (perceived precision). The models outside that class distort whatever is left—how the present is carried forward—and leave a signature in the wedge, a pattern of signs across horizons.

Table 1. Leading models inside the wedge's null

Model	β_{CG}	β_{BGMS}	τ_h	Reference
Block 1: Full-information rational expectations (FIRE)				
FIRE	0	0	0	—
Block 2: Information frictions: rational expectations on any information set				
Information rigidity	$+$	0	0	Coibion and Gorodnichenko (2015)
Rational inattention	$+$	0	0	Kohlhas and Walther (2021)
Imprecise memory	$+$	$-$	0	Sung (2026)
Block 3: Misperceived precision: correct perceived signal–outcome covariance				
Diagnostic Kalman filter	$\pm$	$-$	0	Bordalo et al. (2020)
Ambiguity aversion	$\pm$	$-$	0	Huo, Pedroni and Pei (2024)
Overconfidence	$\pm$	$-$	0	Adam, Kuang and Xie (2025)

Notes: β_{CG} is the coefficient of the consensus forecast errors on the consensus forecast revisions (Coibion and Gorodnichenko 2015); β_{BGMS} is its individual-level counterpart (Bordalo et al. 2020); τ_h is the wedge defined in Definition 1. Each entry is the sign of the population value the diagnostic takes in the model: 0 exactly zero, $+$ positive, $-$ negative, $\pm$ either sign, depending on parameter values. The signs of β_{CG} and β_{BGMS} are model-specific: each is taken in the specification of the paper in the last column. The zero wedges are general results: they hold for every model in the class each block names, with no more assumption needed in the first two blocks, and Assumption 1 needed in the third (see Appendix A2 for more details).

Every interpretation from Table 1 for a nonzero CG or BGMS coefficient sits inside the zero-wedge null. CG read their coefficient as information rigidity, sticky à la Mankiw and Reis (2002) or noisy à la Woodford (2003); BGMS read theirs as diagnostic expectations.[10] Each is a row of Table 1, and Propositions 1 and 2 place the whole class each block names inside the null with it, so a nonzero wedge cannot be any of them.[11]

[10] Complementarily, Kohlhas and Walther (2021) read the pair formed by the CG regression and the error-on-outcome regression in equation (5) as rational inattention asymmetric across states. Sung (2026) reads the BGMS regression as imprecise memory under rational expectations, Huo, Pedroni and Pei (2024) read it as ambiguity aversion, and Adam, Kuang and Xie (2025) read it as overconfidence in private information.

[11] Angeletos, Huo and Sastry (2020) also reach a misperception outside that class, but their approach must assume an AR(1) economy. A nonzero wedge assumes no functional form for the law of motion: it rejects rational expectations for any process, and identifies misperceived propagation under Assumption 1.

1.4 Scope of the wedge

Individual forecasts. The null aggregates. The proofs of Propositions 1 and 2 run forecaster by forecaster, and the two covariances the null equates, $\mathrm{cov}(z_{t+h}, \mathbb{F}_t z_t)$ and $\mathrm{cov}(\mathbb{F}_t z_{t+h}, z_t)$, are means over forecasters of their individual counterparts. If each forecaster's own forecasts carry a zero wedge, so does the consensus; the converse fails, as departures of opposite sign can cancel. The regression of forecast errors on forecast revisions, by contrast, poses distinct hypotheses at the two levels: under information rigidity, the BGMS coefficient is zero for every forecaster while the CG coefficient is positive (Table 1).

The wedge pairs every forecast with an outcome; classical measurement error in either is orthogonal to the other and drops out exactly, at both individual and consensus levels.

Learning models. Assumption 1 fixes the loadings; a forecaster who re-estimates them from realized outcomes falls outside it. In the leading case, constant-gain learning (Evans and Honkapohja 2001; Eusepi and Preston 2011), the one-step wedge τ_1 of a learner who observes the present is proportional to the average drift of her loadings (Proposition A1), an identity that holds exactly when the belief update takes the signal variance as known, and approximately, with an error that shrinks as the gain gets smaller, when the variance needs to be estimated too.

Under a constant gain, beliefs never settle; they keep fluctuating around a stationary distribution, within which the average drift is zero. A learner in that state has no one-step wedge, even if her regression leaves out some relevant observables. Beyond one step, the identity is likewise only approximate, with an error that shrinks as the gain gets smaller (Appendix A3).

A Bayesian learner is rational under her own prior: if the true parameter were a draw from it, Proposition 1 would imply a zero wedge. However, along a sample, her forecasts lean on the prior until the posterior concentrates, and the wedge is tilted toward the prior for as long as learning takes; near a unit root, where persistence is learned slowly (Farmer, Nakamura and Steinsson 2024), a wedge is hard to tell from learning in progress.

Therefore, unsettled learning is a microfoundation of the wedge, not an alternative to it, and the wedge fades as the perceived law of motion converges to its actual counterpart.

Structural change. Assumption 1 also requires covariance stationarity, which makes the wedge a single number at each horizon; under structural changes, the wedge is window-specific, and a pooled sample mixes regimes. When the series are unbounded, the unconditional moments that define the wedge are ill-behaved, so the wedge is least informative there. Section 2 therefore checks the wedge window by window and takes the unbounded series in first differences.

2 The fact: three headline series

This section establishes the headline fact: forecasters under-propagate changes in the unemployment rate and the 3-month Treasury-bill rate but over-propagate the CPI inflation rate. This pattern holds for the consensus and individual forecasts in the Survey of Professional Forecasters, for the Federal Reserve Greenbook, and for the Livingston Survey. Set against the information friction in the same forecast errors, the wedge is of the same size for the two rates. Alternative diagnostics can misread information frictions as sizable misperceived propagation.

2.1 The headline

The headline series are three long-running variables from the U.S. Survey of Professional Forecasters (SPF): the unemployment rate, the 3-month Treasury-bill rate, and CPI inflation, over the forecast origins 1984Q1–2018Q4. A level is a variable in the form models of aggregate fluctuations use—an unemployment, interest or inflation rate—and a change is its first difference. Each headline series enters in the least-differenced form whose realized outcome is not close to a unit root. The two rates, with first-order autocorrelations of 0.98 and 0.99 in levels, therefore enter as quarterly changes, and inflation, at 0.33, as a level; Section 3 reports the other form of each. The sample ends before the pandemic.

Consensus forecasts under-propagate changes in both rates (Figure 1), with average wedges of $+0.24$ (0.10) for the unemployment rate and $+0.21$ (0.06) for the Treasury-bill

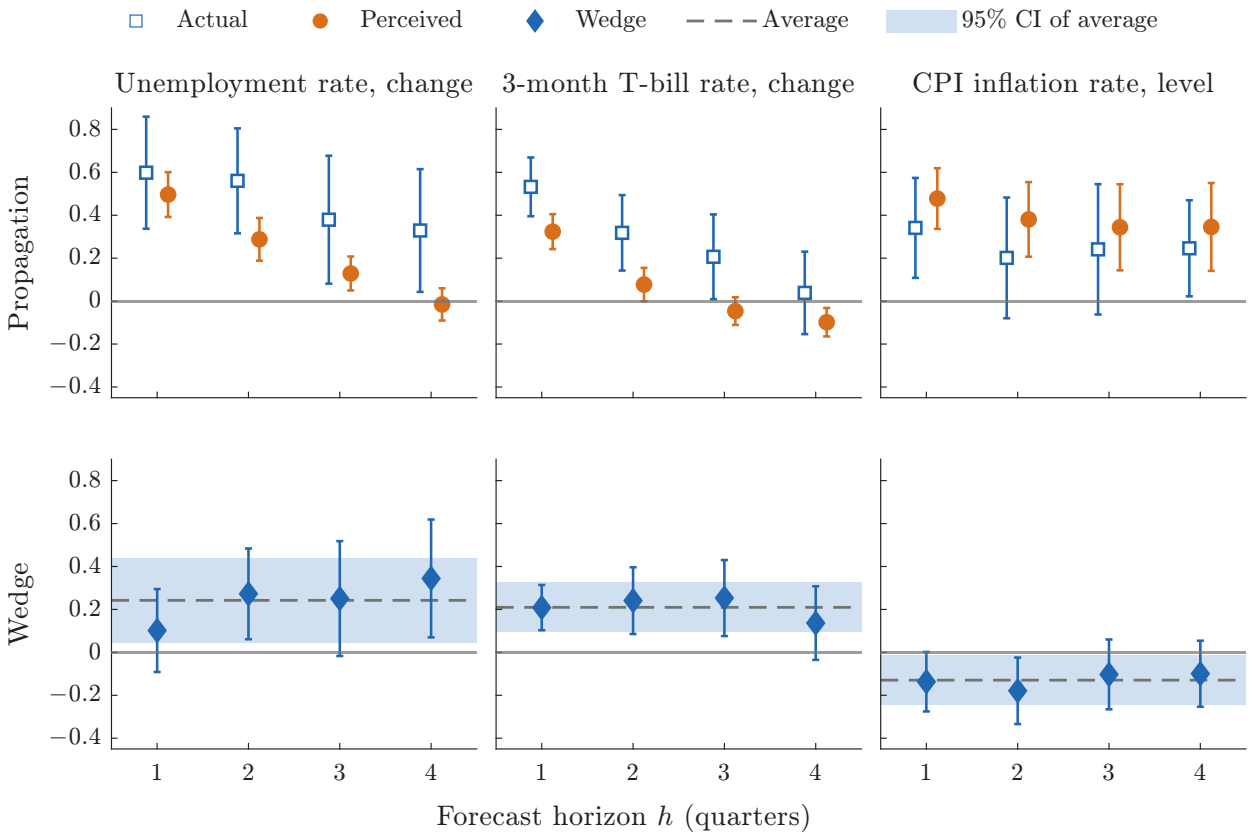


Figure 1. The propagation wedge in the three headline series

Notes: For each of the headline series, the top row plots actual propagation a_h (open squares) and perceived propagation $\widetilde{a}_h$ (filled circles) at horizons $h = 1, \ldots, 4$ quarters, and the bottom row plots the propagation wedge $\tau_h = a_h - \widetilde{a}_h$ (filled diamonds) and its four-horizon average $\overline{\tau}$ (dashed line). A positive wedge is under-propagation and a negative one over-propagation. Forecasts are consensus (mean) forecasts from the Survey of Professional Forecasters made in survey quarters 1984Q1–2018Q4; the unemployment and Treasury-bill rates enter as quarterly changes and CPI inflation as a level (Table 3). Whiskers and the shaded band are 95 percent confidence intervals based on Newey–West standard errors with four lags; Appendix B states the estimator and the construction of each series.
Sources: Survey of Professional Forecasters and Real-Time Data Set for Macroeconomists (second releases), Federal Reserve Bank of Philadelphia; FRED series `TB3MS` (3-month Treasury bill, secondary market).

rate. Averaged across horizons, they carry forward about half of the actual propagation of a change in the unemployment rate ($\tilde{\bar{a}} = 0.22$ against $\bar{a} = 0.47$) and less than a quarter of that of a change in the Treasury-bill rate (0.06 against 0.27). They over-propagate CPI inflation, with an average wedge of -0.13 (0.06), carrying forward about one and a half times its actual propagation (0.39 against 0.26). Every horizon-by-horizon estimate has the sign of its series' average, and so does the average in each of the 81 windows of 60 consecutive forecast origins within 1984–2018 (Appendix E3).

2.2 The forecasters

The fact is not made by averaging: estimated on the individual responses, the wedge matches the consensus estimate to within 0.02 in each of the three series. The individual estimate also settles a condition the published consensus cannot. Section 1.1 carries the null from forecasters to the consensus only when the nowcast and every horizon are averaged over the same forecasters with the same weights at each date; the survey's published mean does not enforce that, and the individual estimate does, because every response it uses reports the nowcast and all four forecasts. The magnitudes are a separate question, and theory leaves it open: a zero wedge for every forecaster makes the consensus wedge zero, but the consensus wedge divides by a covariance that differs from forecaster to forecaster, so it is a ratio of averages rather than an average of the forecasters' own wedges. The second estimate in each panel of Figure 2 answers it in the data: pooling the roughly 220 forecasters who answer for each series over the same origins, with every series centered on the forecaster's own mean, the four-horizon averages are 0.22 (0.11), 0.21 (0.06) and -0.14 (0.06), which differ from the consensus estimates by -0.02, 0.00 and -0.01, with standard errors of 0.02, 0.02 and 0.05 computed jointly with them. The individual wedge carries the consensus sign at all twelve of the horizon-by-series estimates, and the two differ significantly at one: the unemployment rate four quarters ahead, where the individual wedge is smaller by 0.05 (0.02) (Figure A2). One difference in twelve at the 5 percent level is what no difference at all would produce on average. The regression of the forecast error on the forecast revision, the benchmark diagnostic of the survey literature, behaves differently at the two levels, as Section 1.4 anticipates: its coefficient changes sign

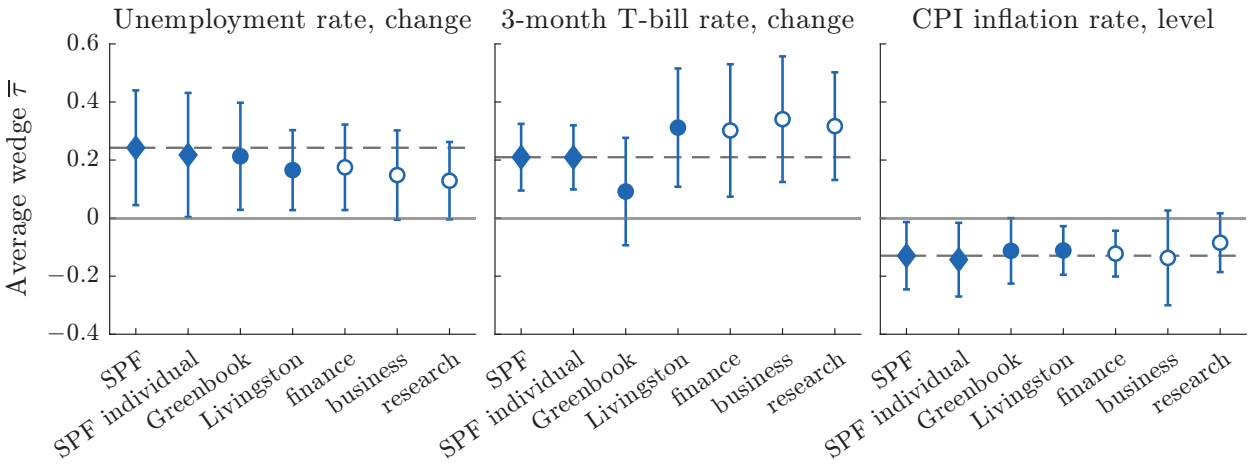


Figure 2. The propagation wedge in individual forecasts and across forecaster populations

Notes: Each marker is the four-horizon average wedge $\overline{\tau}$ in one forecaster population, with its 95 percent confidence interval. Diamonds are the SPF consensus estimate of Figure 1 (also the dashed line) and the estimate pooled over individual SPF responses; filled circles are the Federal Reserve Board staff's Greenbook projections and all Livingston Survey forecasters; open circles are the Livingston affiliation groups. The samples are 1984Q1–2018Q4 for the SPF and the Greenbook, whose Treasury-bill projection ends in 2014Q4, and the semi-annual surveys of 1992–2018 for the Livingston Survey, whose two semester horizons span the same year as the SPF's four quarters. Standard errors are Newey–West with four lags (two for the Livingston Survey) and, for the individual responses, Driscoll–Kraay over forecast origins; Appendix B gives the construction of each population's forecasts.
Sources: Survey of Professional Forecasters, Tealbook (formerly Greenbook) Data Set and Livingston Survey, Federal Reserve Bank of Philadelphia; realized outcomes as in Figure 1.

between individual forecasters and the consensus (Bordalo et al. 2020), while the wedge does not move.

Nor is the fact the SPF's alone. The figure's remaining fifteen estimates come from two other populations: the Federal Reserve Board staff, through its Greenbook projections, and the Livingston Survey, semi-annual and running from 1992, taken as a whole and split by the forecaster's affiliation into finance, business and research. Every one of the fifteen carries the sign of its series. The Greenbook is close to the SPF for the unemployment rate (0.21 against 0.24) and for inflation (-0.11 against -0.13); its Treasury-bill wedge, on the shorter sample that ends in 2014 where the staff's bill projection ends, is positive but smaller and imprecise, 0.09 (0.09). Every Livingston estimate but one lies inside the 95 percent confidence interval of the SPF estimate for the same series; the exception is the business group's Treasury-bill wedge, 0.34 (0.11), which lies above it.

That the Board staff shows the same pattern answers the averaging question a second

way, by a route the individual estimate cannot take: the staff is one institutional forecaster, and no survey average is involved. Its projections are also the ones the Federal Open Market Committee (FOMC) weighs when it sets policy, so the misperception documented here enters the policy process at its input.

2.3 The magnitude

Because τ_h is zero across the entire information-friction class, a nonzero wedge identifies misperceived propagation; this subsection asks how large it is relative to information frictions. The comparison is made on T_h, the covariance of the consensus forecast error at horizon h with the realized current value, scaled by the wedge's denominator. Adding and subtracting the future's covariance with the nowcast, $\operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t)$, splits it exactly into two terms:

$$\underbrace{\frac{\operatorname{cov}(z_{t+h} - \mathbb{F}_t z_{t+h},\, z_t)}{\operatorname{cov}(\mathbb{F}_t z_t,\, z_t)}}_{T_h} = \underbrace{\tau_h}_{\text{misperceived propagation}} + \underbrace{\frac{\operatorname{cov}(z_{t+h},\, z_t - \mathbb{F}_t z_t)}{\operatorname{cov}(\mathbb{F}_t z_t,\, z_t)}}_{r_h \text{ (information friction)}}. \tag{11}$$

The first component is the wedge of Definition 1, which vanishes under correct perception whatever the information set (Proposition 2). The second, r_h, is the future's comovement with the part of the present the nowcast misses, $z_t - \mathbb{F}_t z_t$, and it vanishes under full information whatever the perceived dynamics. Each component is thus identified by the null that switches it off: a nonzero τ_h requires misperceived propagation and a nonzero r_h an imperfect nowcast, whatever else holds. Their sizes are a separate question, taken up at the end of this subsection.

The benchmark measure of information rigidity, due to Coibion and Gorodnichenko (2015), regresses the same forecast error on the forecast revision $\mathbb{F}_t z_{t+h} - \mathbb{F}_{t-1} z_{t+h}$ rather than on the realized present. Its coefficient is zero under full-information rational expectations and responds to information rigidity, but it also responds to misperceived propagation: with full information and a misperceived persistence, the revision predicts the error while $r_h = 0$.[12] The comparison is therefore made on T_h, where the two channels

[12] With $z_{t+1} = \rho z_t + \varepsilon_{t+1}$, full information and forecasts $\mathbb{F}_t z_{t+h} = \widetilde{\rho}^h z_t$, the error is $(\rho^h - \widetilde{\rho}^h) z_t$ plus

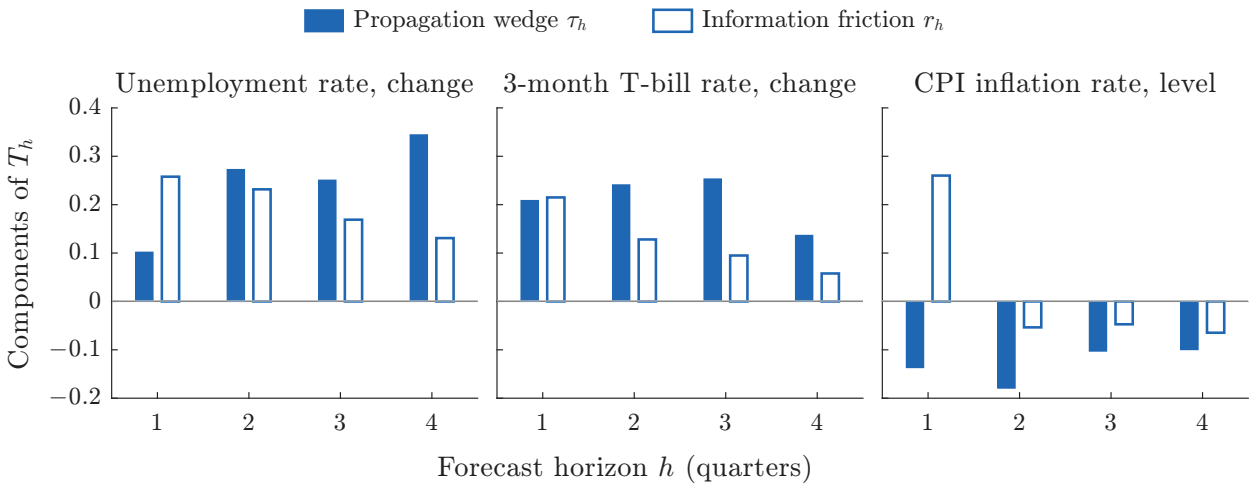


Figure 3. The two components of the predictable forecast error

Notes: Series, transformations, sample and sources are those of Figure 1. Each panel splits $T_h = \mathrm{cov}(z_{t+h} - \mathbb{F}_t z_{t+h}, z_t)/\mathrm{cov}(\mathbb{F}_t z_t, z_t)$, the covariance of the consensus forecast error at horizon h with the realized current value, scaled by the denominator of the wedge, into the two exactly additive components of equation (11): the propagation wedge τ_h of Figure 1 (filled bars) and the information-friction component $r_h = \mathrm{cov}(z_{t+h}, z_t - \mathbb{F}_t z_t)/\mathrm{cov}(\mathbb{F}_t z_t, z_t)$ (open bars). The wedge is zero whenever forecasters perceive propagation correctly, whatever their information; the information-friction component is zero whenever the nowcast equals the realized value, whatever the perceived dynamics. The three terms are estimated from the same sample moments, so $T_h = \tau_h + r_h$ holds exactly in the sample (Appendix B).

have separate nulls. Little of what the revision predicts is lost by conditioning on the present instead. Adding the revision to the regression of the error on z_t raises the R^2 by at most 0.05, and by less than 0.02 for the Treasury bill and inflation (Table A4). Within the sticky- and noisy-information models of Coibion and Gorodnichenko (2015), where perception is correct and the state is an AR(1) with persistence ρ, $\tau_h = 0$ and $r_h = \rho^h(1-\rho^2)\beta$, with β their revision coefficient, so r_h is their information rigidity expressed in the units of the wedge. Outside those models r_h also counts other reasons the nowcast misses the present—for the Treasury bill at the one-quarter horizon, mostly the part of the quarter that has not yet occurred when the survey is taken[13]—which is why r_h is called the information-friction component rather than information rigidity.

future shocks and the revision is $\tilde{\rho}^h(z_t - \tilde{\rho} z_{t-1})$, so their covariance is $(\rho^h - \tilde{\rho}^h)\,\tilde{\rho}^h(1-\tilde{\rho}\rho)\,\mathrm{var}(z_t)$, which is nonzero whenever $\tilde{\rho} \neq \rho$ and $\tilde{\rho} \neq 0$; here $r_h = 0$ and $T_h = \tau_h = \rho^h - \tilde{\rho}^h$.

[13]The Treasury-bill rate is observed daily, and the survey's current-quarter value is a quarterly average. If the rate is a random walk and forecasters are rational given the rates observed up to a fraction θ of the quarter, then $\tau_h = 0$ and the revision coefficient is zero, yet $r_1 = [(1-\theta)^2/2 - (1-\theta)^3/3]/[2/3 - (1-\theta)^3/3]$ and $r_h = 0$ for $h \geq 2$. The variance of the measured Treasury-bill nowcast error, 0.10 of $\mathrm{var}(z_t)$, implies $\theta \approx 0.42$ and hence $r_1 \approx 0.17$, against a measured 0.21.

Figure 3 plots the two components at each horizon. For the two rates they are of the same size. Averaged over the four horizons, the wedge is 0.24 (0.10) against an information-friction component of 0.20 (0.07) for the unemployment rate, and 0.21 (0.06) against 0.12 (0.04) for the Treasury-bill rate.[14] Neither difference is significant: 0.04 with a standard error of 0.07, and 0.09 with 0.07. The wedge's share of the predictable error, $\overline{\tau}/\overline{T}$, is 0.55 (0.08) for the unemployment rate and 0.63 (0.10) for the Treasury-bill rate. For inflation the information-friction component is indistinguishable from zero, 0.02 (0.05): the predictable error is -0.11 (0.09) and the wedge -0.13 (0.06), so the share exceeds one and is better read as this pair. The comparison is within the part of the error the present predicts, which is 1 to 16 percent of its variance (Table A4); the rest is not predictable from the present at all. What is at stake in this margin is how forecasts respond to the present, not the size of the error it leaves. In the point estimates, the horizon profile is common to the three series: the information-friction component is the larger at $h = 1$, and the wedge is the larger from $h = 2$ on. In the forecast errors that the present predicts, misperceived propagation is thus of the same size as information frictions for the two rates, and the whole of a small predictable error for inflation. The comparison is between the components of the errors forecasters actually make, and away from correct perception the wedge's size is not free of what forecasters observe. The channel is a single gap: actual propagation measured against the nowcast, less the same quantity measured against the realized present. It is zero under full information for any law of motion and zero in a univariate state for any information set, so the friction can move the wedge's size only through it (Appendix D1). Averaged over horizons the gap is 0.03 (0.04), -0.04 (0.03) and 0.10 (0.04): distinguishable from zero only for inflation, where the wedge net of it is -0.23 rather than -0.13.

The split is additive by construction and takes no stand on the revision coefficient of Coibion and Gorodnichenko (2015); Section 3.3 fits a model of the misperception on the wedge and reruns the error–revision regression on forecasts net of that model.

[14] Standard errors, in parentheses, are computed as in Figure 1, with each product replaced by its mean over the four horizons.

2.4 The statistic

Four diagnostics can be read off one variance–covariance structure, and only one of them is zero under correct perception whatever forecasters observe. Three natural alternatives suggest themselves—regress the forecast error on the realized present, regress it on the nowcast, or compare the realized autocorrelation with the forecast's—and the cleanest way to see what separates them from the wedge is to read all four off a single variance–covariance structure: that of the present pair $(z_t, \mathbb{F}_t z_t)$—the realized state and its nowcast—and the future pair $(z_{t+h}, \mathbb{F}_t z_{t+h})$—the outcome and its forecast. Collect their four cross-covariances into

$$M_h = \begin{pmatrix} \operatorname{cov}(z_{t+h}, z_t) & \operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t) \\ \operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t) & \operatorname{cov}(\mathbb{F}_t z_{t+h}, \mathbb{F}_t z_t) \end{pmatrix}, \tag{12}$$

its rows the realized outcome and the forecast, its columns the realized present and the nowcast. Each diagnostic differences an actual-propagation entry of M_h against a perceived-propagation one and normalizes, reporting—like the wedge $\tau_h = a_h - \widetilde{a}_h$ (Definition 1)—*actual minus perceived*, so a positive value is under- and a negative one over-propagation. They differ only in which pair of entries they take—and correct perception makes M_h symmetric, so the wedge, which differences the two off-diagonal entries, measures exactly the departure from that symmetry:

$$\text{error on outcome} = \frac{\operatorname{cov}(z_{t+h}, z_t) - \operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t)}{\operatorname{var}(z_t)} \qquad \text{(left column)}$$

$$\text{error on nowcast} = \frac{\operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t) - \operatorname{cov}(\mathbb{F}_t z_{t+h}, \mathbb{F}_t z_t)}{\operatorname{var}(\mathbb{F}_t z_t)} \qquad \text{(right column)}$$

$$\text{persistence comparison} = \frac{\operatorname{cov}(z_{t+h}, z_t)}{\operatorname{var}(z_t)} - \frac{\operatorname{cov}(\mathbb{F}_t z_{t+h}, \mathbb{F}_t z_t)}{\operatorname{var}(\mathbb{F}_t z_t)} \qquad \text{(main diagonal)}$$

$$\text{propagation wedge} = \frac{\operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t) - \operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t)}{\operatorname{cov}(\mathbb{F}_t z_t, z_t)} \equiv \tau_h \qquad \text{(anti-diagonal)}$$

The left-column entry is, up to its normalization, the predictable error T_h that equation (11) decomposes—the same covariance, entering here as a diagnostic in its own right.

The three rivals pair *like-conditioned* entries, and each inherits the friction its conditioning leaves in. Error on outcome regresses both legs on the realized present z_t, so it vanishes only under full information—a noisy nowcast leaves z_t outside the forecaster's information set. Error on nowcast is the textbook efficiency test, both legs on the nowcast $\mathbb{F}_t z_t$: clean for one rational forecaster, not for a consensus. The persistence comparison sets the realized autocorrelation against the forecast's, two slopes whose regressors differ in variance, so beyond an AR(1) a friction opens a gap with no misperception at all (Example 1). The wedge alone takes the anti-diagonal, pairing each future object with the *other* present object; the friction common to both entries cancels, and $\tau_h = 0$ under correct perception for any law of motion (Proposition 2). The distinction bites for the consensus forecast used throughout. A consensus averages rational but dispersedly informed forecasters and so under-reacts to news, the information rigidity of Coibion and Gorodnichenko (2015); that under-reaction loads the forecast error onto *both* columns of M_h, so the error-predictability tests read misperception where there is none, while the anti-diagonal pairing leaves the wedge untouched.

In a rational consensus calibrated to each series, propagation is perceived correctly by construction and information rigidity is the only friction.[15] Averaged over the four horizons, the error-on-outcome and error-on-nowcast diagnostics nonetheless equal $+0.13$ and $+0.11$ for the unemployment rate and $+0.08$ and $+0.08$ for the Treasury-bill rate, 30 to 45 percent of their values in the SPF, and $+0.06$ and $+0.09$ for inflation, whose SPF values are negative. The persistence comparison equals -0.10 for the unemployment rate, whose dynamics depart most from an AR(1), and -0.005 and $+0.01$ for the other two series. The wedge is zero in all three. The population null is proved for any law of motion (Proposition 2); inference is delta-method HAC unless a block bootstrap is stated, and the

[15] For each series the realized z_t is an AR(2), fitted by least squares to the series of Figure 1 over its 140 forecast origins, with coefficients $(0.38,\ 0.31)$ for the unemployment rate, $(0.57,\ 0.02)$ for the Treasury-bill rate and $(0.34,\ -0.03)$ for inflation. The consensus nowcast places weight K on the gap between the present and its one-step prediction μ_t from past values, $\mathbb{F}_t z_t = \mu_t + K(z_t - \mu_t)$, the consensus under-reaction of Coibion and Gorodnichenko (2015), with K set to one minus the ratio of the standard deviation of the measured nowcast error $z_t - \mathbb{F}_t z_t$ to that of the AR(2) innovation (0.40, 0.61 and 0.41). Forecasts iterate the true AR(2) from the nowcast and the realized past, so every forecast is rational given its information and propagation is perceived correctly. Population values are computed on a simulated path of two million quarters.

estimator's own lean in samples of this length is reported in Appendix B1.

The persistence comparison deserves a second look, because on the SPF itself it reports large gaps with the wedge's signs: averaged over the four horizons, $+0.22$ for the unemployment rate, $+0.26$ for the Treasury-bill rate and -0.36 for inflation (Table A5), against wedges of $+0.24$, $+0.21$ and -0.13. How much of each gap is misperception cannot be read off the comparison: a gap appears with no misperception at all in Example 1 and in the two-component environment of Appendix D2, and separating the two takes an estimated model of the dynamics and of the nowcast, every part of which an analyst must choose and defend. The wedge needs none of it. Zero under correct perception for any law of motion (Proposition 2), its advantage is not a larger number but identification: a statistic whose rejection means misperceived propagation and nothing else, delivered without knowing or estimating the data-generating process.

The disagreement is not confined to the three headline series. On every cell of the panel the survey supports (Section 3.4), the three rivals report a value whose block-bootstrap interval excludes zero in 10, 10 and 15 of the 29 cells, against the wedge's 5 (Table A5). In three cells a rival is of the opposite sign to the wedge with an interval that excludes zero: the two core price changes, on samples that start in 2007, and the change in state and local government spending, where the wedge is $+0.19$ and the persistence comparison -0.30. This is the pattern of the calibrated values above in the data: where the nowcast misses the present, the rivals read the miss as misperception, and the wedge does not.

3 The mechanism: one weight on the present

Section 2 established a pattern and its size; this section asks what account of how forecasts are formed can produce the pattern, and how far that account goes. It fits one weight per series on the wedge and then reads the same weight in data the fit never used: the series' other form, the forecast errors, other sample periods, and the other series the survey covers.

3.1 The leading models against the fact

The sign pattern of Section 2 is itself a test of mechanisms. Any account of how forecasts are formed must produce under-propagation for the change in the unemployment rate and for the change in the Treasury-bill rate, and over-propagation for the CPI inflation rate, in the same forecasts. Table 2 collects the leading accounts outside the null of Section 1.2, applies each to the three headline series as they are analyzed—the two rates in change form, inflation in level form—in the AR(1) environment of Example 2, and records the sign it implies.

Three cells are enough because they form a minimal exclusion set. A bias that moves every cell the same way—toward under-propagation, or toward over-propagation—is ruled out by one cell of each sign, and the headline set supplies exactly that: the two rate changes are under-propagated and the inflation level is over-propagated. No larger panel is needed for this step; the panel enters in Section 3.4.

The verdicts follow from the signs. Cognitive discounting under-propagates every series whose own propagation is positive; all three headline series have positive propagation ($\bar{a} = 0.47$, 0.27 and 0.26), so it predicts three positive wedges and fails on inflation. Over-extrapolation, a perceived persistence of the series above its actual one, over-propagates every series, and diagnostic expectations in the variant that leaves the nowcast undistorted do the same; both fail on the two rates. Diagnostic expectations in their canonical form, which tilt the nowcast and the forecasts alike, lie inside the null and predict no wedge at all (Section 1.2). Natural expectations predict a zero wedge at $h = 1$; the Treasury bill's one-step wedge is 0.21 (0.05). Constant-gain learning that has settled predicts wedges near zero at every horizon, and the two rate wedges are well above zero. Simple models carry each component of the economy they track forward at its own persistence and each component they do not track not at all, so when the economy's dynamics are first order, the multivariate counterpart of the table's environment, a series' wedge is zero if it loads only on tracked components and positive otherwise, and inflation's negative wedge fits neither; with higher-order dynamics a tracked component carries the gap of natural expectations, of either sign (Appendix A2).

The weakest verdicts should be named. The cases against cognitive discounting and simple models rest on the inflation cell alone, whose wedge has a Newey–West t-statistic of 2.2 but a 95 percent block-bootstrap interval, $[-0.26, 0.01]$, that contains zero, and a one-sided probability of about 0.07 under the calibrated null of Section 2.4 (Appendix B1). The level form of the other series carries the same information: all ten level-form wedges in the panel are negative (Section 3.4).

What survives is a level anchor, a pull toward the present. A forecast pulled toward the current level carries every level too far and, because the pull cancels when the forecast is differenced, damps every change: the sign follows the form, not the variable, which is the pattern in the table. Its simplest member is projection bias, the tendency to expect current conditions to persist (Loewenstein, O'Donoghue and Rabin 2003); a belief that part of the current level is permanent produces the same reduced form. The anchor is on the present, not on a long-run value. An anchor on a target—the sense in which inflation expectations are said to be anchored—reverts a level toward the target faster than the data do and so under-propagates it; an anchor on the present over-propagates it. The level cells distinguish the two, and all ten carry the sign of an anchor on the present.

The accounts in the table are the ones macroeconomic models use to generate belief dynamics that rational expectations cannot; every one that delivers a sign for these three cells fails at least one of them, and one delivers none—Bayesian learning, whose sign is set by the prior until the posterior converges. The account that survives carries one free parameter per series, the weight on the present. The next subsection fits it on one form of each series and reads it in data the fit never used.

3.2 Projection bias and the level–change matrix

The most parsimonious level anchor is *projection bias*: the forecaster pulls each forecast toward the nowcast just reported, with weight $\lambda \in [0, 1)$ on it (Loewenstein, O'Donoghue and Rabin 2003),

$$\mathbb{F}_t z_{t+h} = \lambda\, \mathbb{F}_t z_t + (1 - \lambda)\, \mathbb{E}_t z_{t+h}, \tag{13}$$

Table 2. Leading models outside the wedge's null, against the three headline wedges

Model	Sign of the wedge τ_h			Against the data
	Unemployment change	Treasury bill change	CPI inflation level	
Measured (Figure 1)	$+0.24$ (0.10)	$+0.21$ (0.06)	-0.13 (0.06)	
Block 1: Fixed rules, with loadings constant over time				
Natural expectations Fuster, Laibson and Mendel (2010)	0 at $h = 1$	0 at $h = 1$	0 at $h = 1$	fails for the Treasury bill
Diagnostic expectations, nowcast undistorted Bordalo, Gennaioli and Shleifer (2018)	$-$	$-$	$-$	fails for the two rates
Over-extrapolation Angeletos, Huo and Sastry (2020)	$-$	$-$	$-$	fails for the two rates
Cognitive discounting Gabaix (2020)	$+$	$+$	$+$	fails for CPI inflation
Simple models Molavi (2025)	0 if tracked; $+$ if not			fails for CPI inflation, tracked or not
Block 2: Learning, with loadings that respond to realized outcomes				
Constant-gain learning, settled Evans and Honkapohja (2001)	≈ 0	≈ 0	≈ 0	fails for the two rates
Bayesian learning Farmer, Nakamura and Steinsson (2024)	$\pm$, by the prior in transition			no sign
Block 3: Level anchor, a pull toward the current level				
Projection bias Loewenstein, O'Donoghue and Rabin (2003)	$+$	$+$	$-$	matches (Section 3.2)

Notes: The three sign columns are the headline series of Figure 1: the change in the unemployment rate, the change in the 3-month Treasury-bill rate and the level of the CPI inflation rate. Signs follow Definition 1: a positive wedge is under-propagation, a negative one over-propagation, $\pm$ is either sign, and ≈ 0 is a wedge near zero, with an error that shrinks as the learning gain gets smaller. The measured row reports the average wedge $\overline{\tau}$ over $h = 1, \ldots, 4$ from Figure 1, with its standard error in parentheses. Each model row applies a leading account outside the null of Section 1.2 to the three series as they are analyzed, in the AR(1) environment of Example 2 or, for simple models, its multivariate counterpart, a first-order economy; signed entries take actual persistence to be positive. Natural expectations fail on the Treasury bill's one-step wedge, 0.21 (0.05); simple models predict a zero wedge for a series that loads only on components they track and a positive one otherwise, and with higher-order dynamics a tracked component carries the gap of natural expectations, of either sign (Appendix A2); and a belief that part of the current level is permanent has the same signs as projection bias. Appendix A2 derives the rows, and Appendix A3 the two learning rows, which fall outside Assumption 1. The verdicts against cognitive discounting and simple models rest on the CPI inflation cell, whose wedge has a Newey–West t-statistic of 2.2 but a block-bootstrap interval that contains zero (Section 3.4).

at every horizon $h \geq 0$, with $\mathbb{E}_t z_{t+h}$ the forecast that would be reported were the present carried forward correctly. The anchor is therefore observed in the survey, directly or from its reported path, rather than inferred as a latent state. Setting $h = 0$ returns $\mathbb{F}_t z_t = \mathbb{E}_t z_t$: the nowcast anchors the forecasts and is itself undistorted, so the bias touches only the forecast leg. The wedge itself requires none of this—its null constrains the perceived signal–outcome covariance at every horizon, whatever the nowcast (Proposition 2)—and it is the projection-bias reading that fixes the $h = 0$ leg as correct and assigns the whole wedge to the horizons beyond it. Because rational forecast errors are orthogonal to date-t information, actual propagation a_h (Definition 1) is unchanged, while the biased forecast lifts perceived propagation toward the present, $\widetilde{a}_h = \lambda + (1 - \lambda)\, a_h$, so the wedge is

$$\tau_h = -\lambda\,(1 - a_h) < 0 : \tag{14}$$

over-propagation, the inflation case, in which forecasters project the current inflation rate.

The sign flips when the analyzed series is the *change* in a projected level. Differencing equation (13) across horizons cancels the common present-anchor, leaving

$$\mathbb{F}_t \Delta z_{t+h} = (1 - \lambda)\, \mathbb{E}_t \Delta z_{t+h}, \tag{15}$$

a pure proportional damping of expected changes. The nowcast of the change stays rational, so perceived propagation is damped to $\widetilde{a}_h = (1 - \lambda)\, a_h$ and the wedge is

$$\tau_h = \lambda\, a_h > 0 : \tag{16}$$

under-propagation, the rate case, in which forecasters project the current *level* of unemployment or the short rate while the wedge is measured on its change.

A single projection bias $\lambda > 0$ thus produces both signs—the switch is a matter of differences versus levels, not of the sign of λ. The weight is not estimated from the errors. Under the maintained model it is the ratio of the average wedge to the average of the factor it multiplies—$\overline{\tau}/\bar{a}$ in change form, $-\overline{\tau}/(1 - \bar{a})$ in level form, the *ratio of averages*: numerator and denominator are averaged over the horizons before dividing—and the

paper fits it once per series, on that series' headline cell, the form in which Figure 1 reports it: 0.52 for the unemployment rate, 0.77 for the Treasury bill and 0.17 for inflation. Every other use of the weight in this section reads it in data the fit did not touch.

Figure 4 shows the six cells, each headline series in both forms, and the pattern is visible before any fitting: every level is over-propagated and the two persistent changes are under-propagated, so the sign follows the form, not the variable. The test is then untargeted, and it is the *level–change test* referred to below. The weight fitted on a series' headline cell predicts the average wedge of its other form with no free parameter: -0.050 against a measured -0.044 for the unemployment level, -0.083 against -0.072 for the Treasury-bill level, and -0.015 against -0.033 for the change in inflation. None is rejected; the block-bootstrap probabilities of the measured values under the implied ones are 0.79, 0.54 and 0.11 (Table A6 extends the test to every series that supports both forms). The three cells do not confirm the anchor equally: for the two rate levels the implied wedge, about -0.05, is of the size of the lean that the estimator carries in a persistent level in samples of this length (Appendix B1), so those cells confirm the sign but cannot separate the anchor from correct perception, while the change in inflation, where the lean is negligible, can.

The weight cannot be read off a regression of forecast errors on forecasts. In change form that regression is the slope of the outcome on the forecast less one, identically, and its regressor is persistent enough that at $T = 140$ it returns a positive weight under rational expectations (Appendix E9). Under the maintained model the split of Section 2.3 is exact: the wedge is the anchor's part of the error the present predicts, and the information-friction component is what remains once the anchor is removed, for any information structure in the null. The anchor then accounts for 0.55 of that predictable error for the unemployment rate and 0.63 for the Treasury bill, and for essentially all of it for inflation—of a part that is itself 1 to 16 percent of the error's variance.

The matrix uses only how forecasts co-move with the present. As a separate check that does not re-estimate the weight from forecast errors, Appendix E7 uses each series' wedge-implied weight to remove the anchor from forecasts at each of the four horizons

and compares the corrected and original squared errors, in sample and in real time (Table A7).

3.3 The anchor and the error–revision regression

The benchmark diagnostic of the survey literature reads the same errors against the forecast revision: Coibion and Gorodnichenko (2015) regress the error of the forecast for quarter $t+h$ on the revision of that forecast between consecutive surveys, and read a positive slope as information rigidity. Under equation (13) the anchor contributes to that slope. Net of the anchor the forecast of a change is $\mathbb{F}_t \Delta z_{t+h}/(1-\lambda)$ at both surveys, so rerunning the regression on the net forecasts gives a slope β_R with $\beta_{\text{obs}} = \beta_R + \lambda b_z$, where b_z is the slope of the realized change on the revision; in level form the net revision also removes λ times the change in the nowcast. The weight is the wedge's, so λb_z is what the anchor alone puts into the coefficient, and β_R is what remains for everything else the null admits, information frictions and overreaction alike; it is not an estimate of rigidity. Figure 5 plots the measured slope and the slope net of the anchor, each with its interval.

For the unemployment rate the measured slopes are 0.89, 1.93 and 1.69 at $h = 1$, 2 and 3, and the anchor alone accounts for 1.17, 1.55 and 1.22 of them; what remains, -0.28 (0.29), 0.38 (0.31) and 0.47 (0.53), is not distinguishable from zero. The same holds in the four-quarter-average specification of Coibion and Gorodnichenko (2015): the anchor accounts for about three quarters of a measured 1.12 (0.41), and the remainder, 0.29, has a bootstrap interval of $[-0.19, 0.71]$. The sign of the anchor's part is mechanical—a positive weight and a positive b_z make it positive—so the content of the check is its size. The size is also an average of two halves that pull in opposite directions: on 1984–2001, where the weight fitted on that half is 0.68, the anchor over-corrects and the remainder is negative at every horizon; on 2001–2018, where it is 0.45, the anchor under-corrects and the remainder is positive and significant at $h = 2$ and 3 (Table A8). This is the drift of the unemployment weight of Section 3.5 seen from the errors. For the Treasury bill the anchor over-corrects at $h = 1$, where the remainder is -0.53 (0.11); in the four-quarter-average specification it accounts for all of a measured 0.75 (0.19), leaving 0.07 with an interval of $[-0.12, 0.26]$. For inflation the regression finds no signal on 1984–2018, so there is noth-

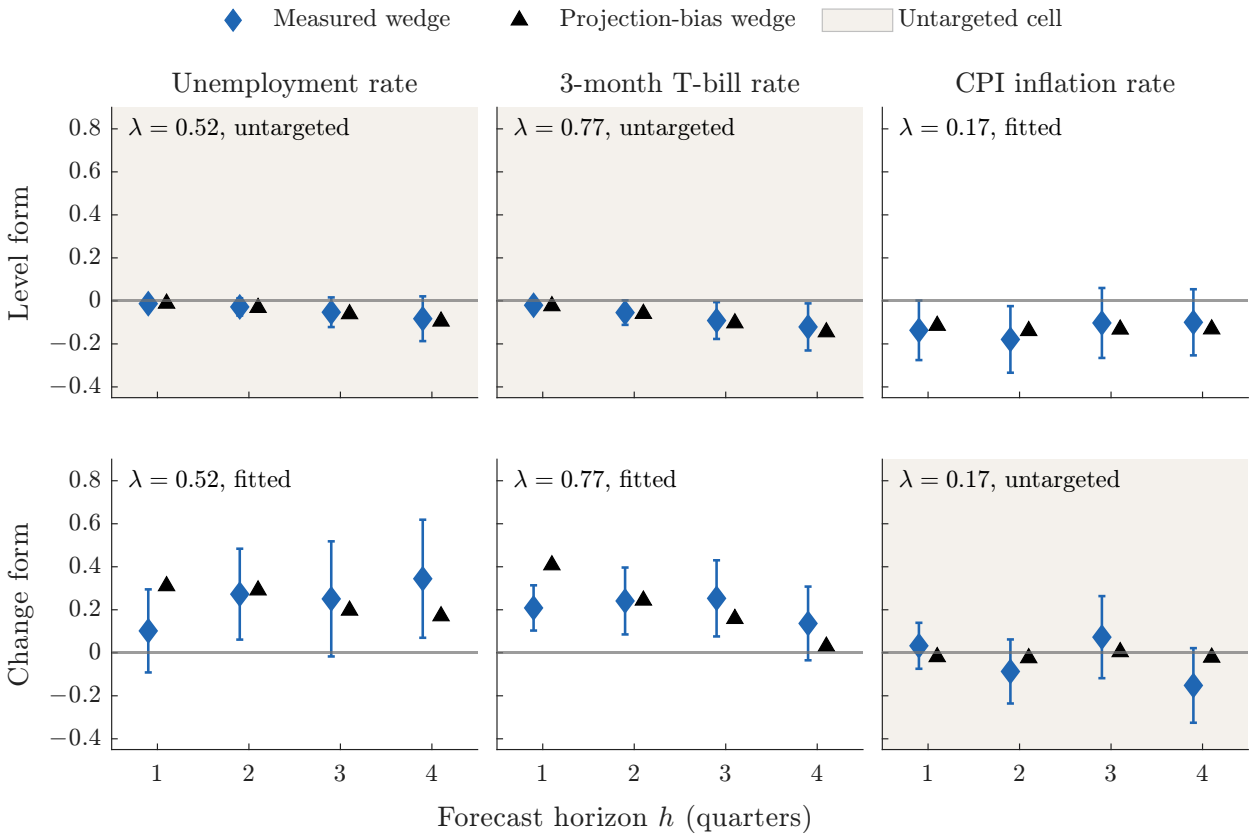


Figure 4. One level anchor fills the level–change matrix

Notes: Series, transformations, sample and sources are those of Figure 1 and, for the untargeted cells, Table 3. Each column is one headline series, in its level form (top row) and its change form (bottom row). Diamonds are the measured wedge τ_h with its 95 percent confidence interval, estimated as in Figure 1. Triangles are the projection-bias wedge, $\tau_h = -\lambda(1 - a_h)$ in level form (equation (14)) and $\tau_h = \lambda a_h$ in change form (equation (16)), each at the cell's own estimated actual propagation a_h. One anchor weight λ per series is fitted on its headline cell alone—the change form for the two rates, the level form for inflation—as the ratio of the average wedge to the average of the factor it multiplies, $\lambda = \overline{\tau}/\bar{a}$ in change form and $\lambda = -\overline{\tau}/(1 - \bar{a})$ in level form, with bars denoting averages over $h = 1, \ldots, 4$. The weight reproduces the headline cell's average wedge exactly and is the 1984–2018 weight of Figure 7. Shaded panels are the untargeted cells, whose triangles carry no free parameter. The vertical scale is that of Figure 1.

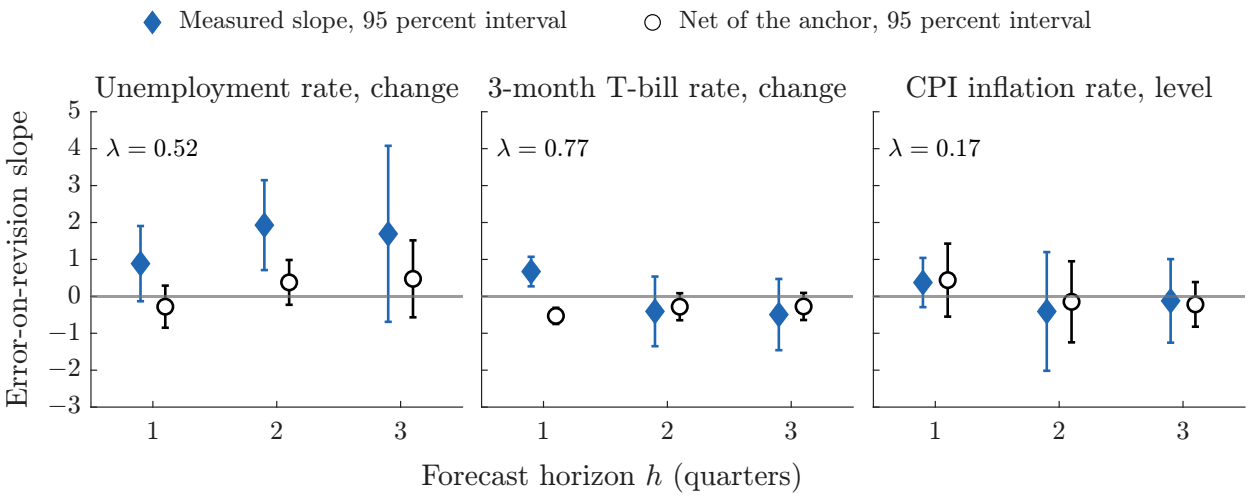


Figure 5. The anchor and the error–revision regression

Notes: Series, transformations, sample and sources are those of Figure 1. Each panel is one headline series, at horizons $h = 1, 2, 3$. Diamonds are the slope of the consensus forecast error on the revision of the same forecast between consecutive surveys, $z_{t+h} - \mathbb{F}_t z_{t+h}$ on $\mathbb{F}_t z_{t+h} - \mathbb{F}_{t-1} z_{t+h}$ (Coibion and Gorodnichenko 2015). Open circles are the same slope on forecasts net of the anchor, which divide every forecast of a future change by $1 - \lambda$—in level form, $(\mathbb{F}_t z_{t+h} - \lambda\, \mathbb{F}_t z_t)/(1 - \lambda)$—at both surveys, with the weight λ of Figure 4. The measured slope less the net slope is the part the anchor alone implies, λb_z in change form, where b_z is the slope of the realized change on the revision; its sign is mechanical for a positive weight. Whiskers are 95 percent confidence intervals based on Newey–West standard errors with four lags; for the net slopes they are computed by the delta method on the stacked moments of the regression and of λ, so they charge the weight's sampling error. The nowcast, $h = 0$, is omitted because in change form its target contains last quarter's level, which is data at the survey; $h = 4$ would need a five-quarter forecast, which the survey does not report. Block-bootstrap intervals, the two halves of the sample and the four-quarter-average specification of Coibion and Gorodnichenko (2015) are in Table A8.

ing to match; in a replication of the benchmark on the GDP deflator the coefficient is 1.31 on 1969–1983 and 0.20 on 1984–2014 (Table A9). The exercise completes the benchmark regression under a stated model; it does not decompose its coefficient.

3.4 Beyond the three series

Section 2.2 showed that the fact holds across forecasters; this subsection carries it, and the mechanism fitted on the three headline series, across every series the survey asks, and Section 3.5 then carries it across time. The level–change test of Section 3.2 needs both forms of a series, so it extends naturally to every series that supports both forms, and the one weight it fits can then be carried to every cell of the panel. Table 3 lists the panel: twenty series in the forms they support, 29 cells in all, each on the longest sample without a missing quarter and every one ending in 2018Q4. The wedge is measured in every cell

as in Figure 1.

Jointly, the 29 wedges are not zero: a moving-block bootstrap over calendar quarters, applied to all cells at once, rejects at $p = 0.01$, and the 19 change-form cells alone at $p = 0.004$; the 10 level-form cells alone do not reject ($p = 0.26$), yet all ten are negative, a count that overstates its evidence, because in samples of this length the estimator leans negative in a persistent level (Appendix B1; Appendix E4 plots every cell with its interval). Cell by cell, 9 of the 29 wedges are significant with Newey–West standard errors and 5 under the bootstrap; the small counts in the level cells are what the mechanism predicts, since the level-form factor $1 - a_h$ is small for a persistent series. A second joint statistic does not reject: the sum of the 29 squared t-statistics is 86 against a bootstrap 95th percentile of 143 ($p = 0.14$). It discards the covariance across cells that the Wald statistic uses, and with five individually significant cells among twenty-nine it has little power against the alternative the mechanism implies, in which most cells carry a small wedge of a predictable sign.

The level–change test extends to the nine series that support both forms, the series the survey reports as a rate (Table 3). For each, the weight fitted on the headline form predicts the sign of the other form's average wedge, and the measured sign matches in all nine; no implied average wedge is rejected at the 5 percent level, PCE inflation coming closest ($p = 0.06$; Table A6). Where the weight is identified, five series, it is higher for the rates—0.52, 0.77 and 1.09 for the unemployment rate, the Treasury bill and the Baa yield—than for inflation, 0.17 for the CPI and 0.20 for the PCE deflator.

The ordering is a result. The three headline weights are not one number: the hypothesis that they are equal is rejected on the same moving-block resamples as the rest of this section (bootstrap $p = 0.010$), and what is rejected is the contrast between the rates and inflation—0.61 [0.42, 0.77] fitted on the two rate cells against 0.17 [−0.02, 0.32] on the inflation cell, a difference of 0.44 [0.20, 0.68] ($p = 0.006$)—while the two rates do not differ from each other ($p = 0.085$). The Baa weight exceeds one, the bound a projection weight should respect, and its interval contains one; it is fitted on thirty-six quarters, and Appendix E6 adds it and the PCE deflator to the test.

Figure 6 carries one weight to every cell. Pooling the three headline cells by the ratio

Table 3. The SPF series: what the survey asks, what the paper builds, and over which sample

Series	SPF code	SPF reports x as	Analyzed as: level	Analyzed as: change	Sample start (all end 2018Q4)
Rates					
Unemployment[a]	UNEMP	rate	x	Δx	1984Q1
3-month T-bill[c]	TBILL	rate	$\ln(1+x\%/4)$	$\Delta\ln(1+x\%/4)$	1984Q1
10-year T-note[c]	TBOND	rate	$\ln(1+x\%/4)$	$\Delta\ln(1+x\%/4)$	1992Q1
Aaa corporate bond[c]	BOND	rate	$\ln(1+x\%/4)$	$\Delta\ln(1+x\%/4)$	1984Q1
Baa corporate bond[c]	BAABOND	rate	$\ln(1+x\%/4)$	$\Delta\ln(1+x\%/4)$	2010Q1
Inflation					
CPI[a]	CPI	rate	$\ln(1+x\%)/4$	$\Delta\ln(1+x\%)/4$	1984Q1
PCE[a]	PCE	rate	$\ln(1+x\%)/4$	$\Delta\ln(1+x\%)/4$	2007Q1
Core CPI[d]	CORECPI	rate	$\ln(1+x\%)/4$	$\Delta\ln(1+x\%)/4$	2007Q1
Core PCE[a]	COREPCE	rate	$\ln(1+x\%)/4$	$\Delta\ln(1+x\%)/4$	2007Q1
GDP deflator[b]	PGDP	index	$\Delta\ln x$	—	1984Q1
Activity					
Real GDP[b]	RGDP	level	—	$\Delta\ln x$	1984Q1
Nominal GDP[b]	NGDP	level	—	$\Delta\ln x$	1984Q1
Real consumption[b]	RCONSUM	level	—	$\Delta\ln x$	1984Q1
Real nonres. investment[b]	RNRESIN	level	—	$\Delta\ln x$	1984Q1
Real resid. investment[b]	RRESINV	level	—	$\Delta\ln x$	1984Q1
Real federal gov.[b]	RFEDGOV	level	—	$\Delta\ln x$	1984Q1
Real state & local gov.[b]	RSLGOV	level	—	$\Delta\ln x$	1984Q1
Industrial production[a]	INDPROD	level	—	$\Delta\ln x$	1984Q1
Housing starts[a]	HOUSING	level	—	$\Delta\ln x$	1984Q1
Payroll employment[a]	EMP	level	—	$\Delta\ln x$	2003Q4

Notes: Forecasts are consensus (mean) forecasts from the Survey of Professional Forecasters (SPF), and x is each series as the survey reports it. The estimator uses x where the survey gives a rate and its log growth where the survey gives an index or a level, which trends. Inflation and interest rates enter as quarterly log rates, together with their first differences: an inflation rate, which the SPF annualizes by compounding, as $\ln(1+x\%)/4$, the quarterly log change in the price index; an interest rate, quoted at a simple annual rate, as $\ln(1+x\%/4)$, the quarterly log return on simple interest at that rate. The unemployment rate enters as reported, being a share rather than a return. A dash marks a form the series cannot support. The estimator works in annualized percent, 400 times these quarterly log rates; a common rescaling leaves the wedge, a ratio of covariances, unchanged. Each sample starts at the first survey quarter at which the nowcast, all four forecasts, the realized current value and its four-quarter lead are present, and runs without a missing quarter to 2018Q4, so that the realizations end in 2019Q4, before COVID-19. Series in bold support both forms over the full window and are the headline series. The change form of the two corporate yields takes the nowcast's lag from the realized yield, because the SPF publishes no previous-quarter value for them; on the 10-year note, where it does, the substitution moves no statistic by more than 0.008.
Sources: Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia. Realized outcomes are second releases: [a] Real-Time Data Set for Macroeconomists, Federal Reserve Bank of Philadelphia, vintages of monthly data; [b] the same, vintages of quarterly data, differenced within the second vintage that contains each quarter; [c] FRED quarterly averages, never revised: the 3-month bill on the survey's secondary-market basis (TB3MS), the 10-year note, and Moody's seasoned Aaa and Baa yields; [d] FRED latest vintage, core CPI having no real-time vintage.

of averages of Figure 4 gives $\lambda = 0.39$, and the other 26 cells are predicted with no free parameter from their own actual propagation. The implied sign, which needs only a positive weight, matches the measured sign in 22 of the 26 held-out cells, which has probability 0.018 under no wedge. The weight the held-out cells themselves want is 0.33, which has probability 0.03 under no wedge once the estimation error shared by the implied and the measured wedges is charged, and below 0.001 with each cell weighted by its precision; it lies 0.06 below the headline weight, with an interval of $[-0.23, 0.14]$ for the difference. The magnitudes go with the signs. In the sample, the wedges implied by the 0.39 weight fitted on the three headline cells account for 0.57 of the sum of squared measured wedges across the 26 held-out cells. What separates one cell's wedge from another's is therefore mostly the cell's own actual propagation, which is what the mechanism says it should be: the weight is common across cells and the factor it multiplies is not.

One weight across series is therefore a summary of two. Asked what they want on their own, the held-out cells of the rate series answer 0.71 $[0.49, 0.85]$, those of the price series 0.38 $[-0.05, 0.63]$ and those of the activity series 0.19 $[-0.13, 0.49]$; the pooled 0.39 sits between, fits the price and activity cells about as well as a weight of their own would, and under-fits the rates, where their own weight raises the held-out R^2 from 0.73 to 0.92. The split is less sharp in the held-out cells than in the fitted ones—their rate-minus-price contrast is 0.33 $[-0.05, 0.83]$, $p = 0.10$—and one weight and two match the same 22 of 26 signs.

3.5 The weight across sample windows

The last dimension the fit did not touch is time.

By equations (14) and (16), the wedge is the anchor weight λ times a factor fixed by actual propagation alone—a_h in change form, $-(1 - a_h)$ in level form—whatever forecasters observe. Actual propagation is a property of the data and differs from one sample to another, so the wedge can move because the factor moves, because the weight moves, or both; the windows separate the two. Figure 7 checks this on every window of 60 consecutive forecast origins, through the last origin whose four-quarter outcome has been released. Sixty is close to the shortest window the implied weight can be read on: it di-

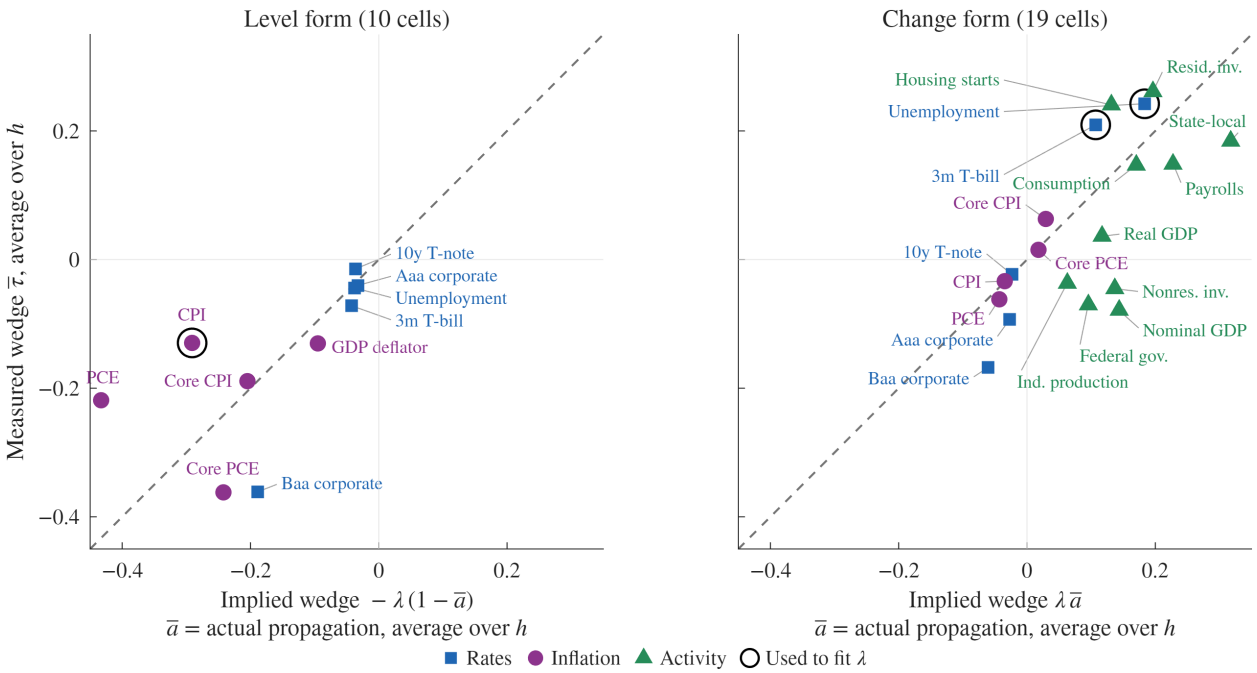


Figure 6. One anchor weight against every cell of the SPF panel

Notes: Series, transformations and samples are those of Table 3, and every wedge is estimated as in Figure 1. Each marker is one cell of the panel—one series in one form—and plots its measured average wedge $\bar{\tau}$ against the wedge one projection-bias weight implies from the cell's own average actual propagation $\bar{a}$, $-\lambda(1-\bar{a})$ in level form and $\lambda\bar{a}$ in change form. The dashed line is the 45-degree line, on which the implied wedge equals the measured one. The weight, $\lambda = 0.39$, is fitted on the three ringed headline cells together by the ratio of averages of Figure 4, taken over the three cells as well as the four horizons: the sum of the two rate-change wedges less the inflation-level wedge, over the sum of the factors they multiply, $\bar{a}$ for the rate changes and $1-\bar{a}$ for the inflation level. It is the estimator of Figure 4 with the three cells added to its sums, so it reproduces that sum of wedges exactly and gives the same weight whether the twelve horizon-by-cell wedges or the three averages are pooled; least squares, which weights each cell by the square of its factor, does neither. The other 26 cells carry no free parameter. The held-out weight of the text is fitted the same way on those 26 cells alone, each wedge multiplied by the sign the anchor predicts for its cell (negative in level form, that of $\bar{a}$ in change form) and each factor taken in absolute value. One weight across series is a summary of two: the headline series' own weights, 0.52, 0.77 and 0.17 (Figure 4), differ (jointly, bootstrap $p = 0.010$; Appendix E6), so the ringed cells lie off the line, and the three untargeted cells of Figure 4 appear here at the common weight rather than at their own series' weights. Probabilities in the text are from a moving-block bootstrap over calendar quarters (8-quarter blocks, 1,999 draws) applied to all cells at once, and Figure A3 plots the measured wedges with their confidence intervals.

vides by the window's own actual propagation, and over five years that divisor wanders near zero. The wedge itself needs far less, and Table A3 separates the two. Its top row plots the average wedge $\overline{\tau}$ with the average of the factor over the same four horizons, $\bar{a}$ for the two rate changes and $-(1-\bar{a})$ for the inflation level; its bottom row plots their ratio, the implied λ.

The three series divide the work between the two factors. For the Treasury bill, the wedge and the weight both hold: the wedge stays positive in every window, between 0.13 and 0.32 both before 2018Q4 and after it, and the implied λ stays between 0.66 and 0.93 within 1984–2018 and between 0.53 and 0.82 after, around its 1984–2018 value of 0.77. For inflation, actual propagation drifts: it falls from 0.43 on average in the windows that end in 1999–2001 to -0.04 in those that end in 2011–2018. Perceived propagation falls with it, from 0.52 to 0.11, so the wedge moves only from -0.09 to -0.15, where a perceived propagation held at its earlier value would have taken it to -0.56; the implied λ is 0.15 and 0.14 in the two groups of windows, and between 0.08 and 0.30 in every window within 1984–2018. For unemployment, actual propagation changes sign: the jump and reversal of 2020 make it negative, about -0.12, in every window that ends in 2020Q3 or later, because they dominate the second moments of any sample that contains them—the seven origins from 2019Q2 to 2020Q4, 4 percent of the extended sample, carry 93 percent of the covariance in the wedge's denominator. The wedge changes sign with it, to about -0.04, where a perceived propagation held at its value in the window that ends in 2018Q4 would have taken it to -0.40; the implied λ keeps its sign, between 0.31 and 0.34.

Within 1984–2018 the weight moves little: across the 81 windows its standard deviation is 0.14, 0.06 and 0.05 for the unemployment rate, the Treasury bill and inflation. The windows describe the weight; they do not test whether it is constant, since consecutive windows share fifty-nine of their sixty quarters and their 81 estimates are far from independent.

Run back further than the figure does, as far as each survey question reaches, all three drift, and they drift the same way (Figure A1). Between the first third of a series' windows and the last, the weight falls from 0.65 to 0.40 for the unemployment rate, from 0.84 to 0.68

for the Treasury bill and from 0.19 to 0.06 for inflation. The fall is about twice as large for the unemployment rate as for the other two. A common slow decline in the weight, over forty years and in all three series, is what these windows show; it is not a feature any one series' environment can account for.

Whether what remains of the movement answers to the environment is a question this paper asks and does not settle. Coibion and Gorodnichenko (2015) settle the corresponding question for information rigidity on the Great Moderation, and the survey allows the same test for one series: it has asked for the unemployment rate since 1968Q4, and all but five of its rounds carry the whole fan of forecasts, so its windows reach well behind the 1984Q1 break of McConnell and Perez-Quirós (2000). The Treasury-bill and CPI questions begin only in 1981Q3 and reach the break by ten quarters. For the unemployment rate the sign is the benchmark's, and the strength of it depends entirely on how far the windows are allowed to run. Inside the estimation sample the weight and the volatility of the series have a correlation of -0.94; across all 159 windows through 2024Q2 it is -0.63, and across the 142 that do not contain the pandemic it is -0.37. Sorted by volatility rather than by date, those 159 windows give weights of 0.68, 0.51 and 0.52 from the least to the most volatile third: the quiet third carries the higher weight and the other two are the same, so the relation is a contrast between calm and everything else rather than a gradient. A regression that uses no ratio gives a weight lower by 0.15 per standard deviation of volatility over all the windows and by 0.08 once the pandemic ones are set aside. For the other two series the correlations with volatility, $+0.49$ and -0.28, are values the constant-weight null of Appendix E3 produces 34 and 48 percent of the time. Inflation's weight instead falls with the level of inflation in the windows that reach past 2018Q4, which all contain one episode. The three series therefore give three answers, and what feature of the environment the anchor answers to is not settled here; Figure A1 plots both candidate features against the weight for all three series and reports the tests.

Whatever part of the movement is real belongs to the anchor: the wedge is zero under any information friction (Proposition 2), so no friction can move the weight.

The weight can also be fitted in real time and applied. Netting out the anchor at the

weight re-fitted at each origin lowers the mean squared error of the unemployment level forecasts by about a sixth, in sample and out of sample, more than the rigidity correction of Coibion and Gorodnichenko (2015) does at the same origins, though neither gain is significant. No other cell has a material or significant pooled gain—the largest, CPI inflation in change form, is 1.3 percent in sample and 2.1 percent in real time ($p = 0.14$)—and for the Treasury bill netting out the anchor raises the error (Table A7).

For a model of belief formation the lesson is a parameter: one weight on the present, about 0.3 to 0.4 when pooled, reproduces the sign of the wedge across the survey, and where it is identified series by series it is higher for rates than for inflation. Over the sample windows it drifts, downward in all three series and most for the unemployment rate.

4 Conclusion

A forecast can go wrong in what forecasters observe of the present or in how they carry the present into the future, and this paper identifies the second. The propagation wedge is zero for every forecaster who carries whatever she knows about the present forward by the economy's own dynamics, however noisily she observes the present and however slowly she updates, so a nonzero wedge rejects rational expectations without a model of the economy or of the information forecasters hold, and under linear forecasting rules it identifies misperceived propagation. Its null contains the nulls of the two classic survey diagnostics, so a rejection of it rules out the most. In the Survey of Professional Forecasters the wedge is not zero. Forecasts carry forward about half of the actual propagation of a change in the unemployment rate, less than a quarter of that of a change in the Treasury-bill rate, and about one and a half times the actual propagation of CPI inflation. The individual responses carry the sign as the consensus does, the Federal Reserve Board staff and the Livingston Survey share it, and it holds in every fifteen-year window of the sample.

The findings bear on how expectations are measured and modeled. First, carrying the present forward is a margin of expectation error in its own right. In the part of the



Figure 7. The wedge and the implied anchor weight on rolling sample windows

Notes: Series, transformations and sources are those of Figure 1; each window holds 60 consecutive forecast origins and is plotted at its last origin. The top row plots the average propagation wedge $\overline{\tau}$ over $h = 1, \ldots, 4$ (blue) with its 95 percent confidence band, and the average over the same horizons of the factor that multiplies the anchor weight in the projection-bias wedge (orange): $\bar{a}$ for the two rate changes (equation (16)) and $-(1-\bar{a})$ for the inflation level (equation (14)). The bottom row plots the implied anchor weight, the ratio of the two lines above—$\lambda = \overline{\tau}/\bar{a}$ in change form, $\lambda = -\overline{\tau}/(1-\bar{a})$ in level form—with its 95 percent confidence band, and its 1984–2018 value (dashed). The 81 windows to the left of the dotted line lie within the estimation sample, their last origins running from 1998Q4 to 2018Q4. The windows on the gray ground end after 2018Q4; those ending in 2020Q1 or later contain the 2020 quarters among their origins, and those ending in 2019 reach them only through their four-quarter outcomes; their last origins run to 2024Q2 for unemployment, 2025Q3 for the Treasury bill and 2024Q3 for inflation, the last origins whose four-quarter outcome has been released. The estimator of Figure 1 is applied window by window, each series centered on its own mean within the window, and the bands are delta-method with Newey–West standard errors with four lags. In change form the window ratio is biased toward zero and its bands are narrower than its sampling spread, so the text compares window values only with window values; Appendix E3 gives the bias, the constant-weight null and the window-level formulas.

forecast error that the present predicts, the wedge is as large as the part that vanishes under full information for the two rates; for inflation that predictable part is small, and the wedge is the whole of it. That predictable part is one to sixteen percent of the error's variance, so the margin governs how forecasts respond to the present rather than how large their errors are. Second, the misperception has a shape that a model of expectations must reproduce. Its sign follows the form of the series rather than the variable: a level is carried too far and a persistent change too little. No bias that moves every forecast the same way produces that pattern; a level anchor, a pull toward the present whose simplest form is projection bias, does. One weight on the present per series reproduces both signs and predicts the sign of the wedge in the series' other form with no free parameter. Third, the benchmark regression of the survey literature, of the forecast error on the forecast revision, conditions on what forecasters learn about the present and does not separate the two margins. Under the maintained model the anchor contributes to its slope, and for the unemployment rate the slope net of the anchor is not distinguishable from zero at any horizon, though with wide intervals. Finally, the Federal Reserve staff's projections carry the same signs, so the misperception enters the policy process at its input.

The wedge also opens questions this paper does not answer. The first is how movements transmit across series. Because the anchor is a property of each series' forecast, it predicts with no further parameter how the forecast of one series comoves with the present of another. A departure from that prediction would mean that forecasters misperceive how a movement in unemployment carries into inflation or the policy rate, not only how each series persists, and the survey's forecasters report all three. The second is the weight on the present itself. It is about 0.6 for the two rates and about 0.2 for inflation, it drifts downward over forty years in all three series, and what sets it is open. The third is whether the misperception concentrates on particular disturbances. Conditioning the wedge on identified shocks instead of the nowcast does not answer it: a forecaster who cannot tell which shock hit carries its effect forward by an average of the dynamics of all of them, even when she perceives each correctly, so a shock-by-shock wedge needs a null of its own. The fourth is what the anchor does to decisions. A path pulled toward the quarter in which it is written moves only a fraction of the way the correctly perceived

path moves. How that fraction enters the decisions of households, firms and the central bank—through the slope of an expectations-augmented Phillips curve, the transmission of forward guidance, or a policy rule that responds to staff forecasts and so partly to the present—is a question for a model of decisions rather than of forecasts. The statistic needs only a nowcast and forecasts from the same respondents, so the same questions can be put to the expectations of firms and households wherever a survey elicits both. Models of expectations are now asked whether they reproduce the information-rigidity coefficient; they can also be asked whether the propagation they imply is the propagation forecasters use.

References

Adam, Klaus, Pei Kuang, and Shihan Xie. 2025. "Overconfidence in Private Information Explains Biases in Professional Forecasts." *Journal of Monetary Economics*, 155: 103839.

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry. 2020. "Imperfect Macroeconomic Expectations: Evidence and Theory." In *NBER Macroeconomics Annual 2020, vol. 35*, ed. Martin Eichenbaum, Erik Hurst and Jonathan A. Parker, 1–86. Chicago: University of Chicago Press.

Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo. 2026. "Smooth Diagnostic Expectations." *Review of Economic Studies*. Forthcoming.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2018. "Diagnostic Expectations and Credit Cycles." *Journal of Finance*, 73(1): 199–227.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review*, 110(9): 2748–2782.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105(8): 2644–2678.

da Silveira, Rava Azeredo, Yeji Sung, and Michael Woodford. 2024. "Optimally Imprecise Memory and Biased Forecasts." *American Economic Review*, 114(10): 3075–3118.

Eusepi, Stefano, and Bruce Preston. 2011. "Expectations, Learning, and Business Cycle Fluctuations." *American Economic Review*, 101(6): 2844–2872.

Evans, George W., and Seppo Honkapohja. 2001. *Learning and Expectations in Macroeconomics.* Princeton:Princeton University Press.

Farmer, Leland E., Emi Nakamura, and Jón Steinsson. 2024. "Learning About the Long Run." *Journal of Political Economy*, 132(10): 3334–3377.

Fuster, Andreas, David Laibson, and Brock Mendel. 2010. "Natural Expectations and Macroeconomic Fluctuations." *Journal of Economic Perspectives*, 24(4): 67–84.

Gabaix, Xavier. 2019. "Behavioral Inattention." In *Handbook of Behavioral Economics: Applications and Foundations 1*, Vol. 2, ed. B. Douglas Bernheim, Stefano DellaVigna and David Laibson, 261–343. Amsterdam: North-Holland.

Gabaix, Xavier. 2020. "A Behavioral New Keynesian Model." *American Economic Review*, 110(8): 2271–2327.

Goldstein, Nathan, and Yuriy Gorodnichenko. 2023. "Expectations Formation and Forward Information." *NBER Working Paper 29711*.

Huo, Zhen, Marcelo Pedroni, and Guangyu Pei. 2024. "Bias and Sensitivity under Ambiguity." *American Economic Review*, 114(12): 4091–4133.

Keane, Michael P., and David E. Runkle. 1990. "Testing the Rationality of Price Forecasts: New Evidence from Panel Data." *American Economic Review*, 80(4): 714–735.

Kohlhas, Alexandre N., and Ansgar Walther. 2021. "Asymmetric Attention." *American Economic Review*, 111(9): 2879–2925.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics*, 118(4): 1209–1248.

Lovell, Michael C. 1986. "Tests of the Rational Expectations Hypothesis." *American Economic Review*, 76(1): 110–124.

Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review*, 99(3): 769–803.

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics*, 117(4): 1295–1328.

McConnell, Margaret M., and Gabriel Perez-Quirós. 2000. "Output Fluctuations in the United States: What Has Changed since the Early 1980s?" *American Economic Review*, 90(5): 1464–1476.

Mincer, Jacob, and Victor Zarnowitz. 1969. "The Evaluation of Economic Forecasts." In *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, ed. Jacob Mincer, 3–46. New York: National Bureau of Economic Research.

Molavi, Pooya. 2025. "Simple Models and Biased Forecasts." *Working Paper*.

Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica*, 29(3): 315–335.

Reis, Ricardo. 2020. "Comment on 'Imperfect Expectations: Theory and Evidence'." In *NBER Macroeconomics Annual 2020, vol. 35*, ed. Martin Eichenbaum, Erik Hurst and Jonathan A. Parker, 99–111. Chicago: University of Chicago Press.

Sung, Yeji. 2026. "Macroeconomic Expectations and Cognitive Noise." *American Economic Review, conditionally accepted*.

Woodford, Michael. 2003. "Imperfect Common Knowledge." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 25–58. Princeton University Press.

Appendix A. Supporting calculations for Section 1

A1. The diagnostic Kalman representation

In Example 4, the steady-state rational nowcast satisfies $\mathbb{E}[z_{t+h}|\mathcal{I}_{i,t}] = \rho^h x_{i,t}$ for every h; combined with the law of iterated expectations, this delivers the true second moments of the signal vector $\mathbf{S}_{i,t} = (x_{i,t}, x_{i,t-1})^\top$,

$$\boldsymbol{\Sigma}_i = V \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \qquad \mathbf{C}_{i,h} = V\rho^h \begin{pmatrix} 1 \\ \rho \end{pmatrix}. \tag{A1}$$

Direct multiplication verifies that $\widetilde{\boldsymbol{\Sigma}}_i \boldsymbol{\beta}_{i,h} = \mathbf{C}_{i,h}$ for the perceived variance in equation (10), so those perceived moments reproduce the diagnostic loadings $\boldsymbol{\beta}_{i,h} = \rho^h(1+\theta, -\theta\rho)^\top$ exactly. The determinant of equation (10) is $V^2(1-\rho^2)/(1+\theta)$, and its leading entry is positive, so the matrix is positive definite for $\theta > -1$ and $|\rho| < 1$. When $\theta > 0$, the leading entry falls below V: the forecaster behaves as if her current state estimate were more precise than it is, while perceiving its propagation correctly.

A2. Model placements in Tables 1 and 2

This subsection derives every entry of the two tables, row by row. Two tools do most of the work. In the univariate environment of Example 2, a forecaster who carries a nowcast formed from date-t information forward at the true persistence, $\mathbb{F}_t z_{t+h} = \rho^h \mathbb{F}_t z_t$, has $\widetilde{\rho} = \rho$ and hence $\tau_h = 0$ at every horizon, however the nowcast is formed; one who carries it forward at a single perceived persistence $\widetilde{\rho} \neq \rho$ has $\tau_h = \rho^h - \widetilde{\rho}^h$, whose sign is that of $\rho - \widetilde{\rho}$ when $\rho > 0$. Beyond that environment, Part (ii) of Proposition 2 places inside the null every linear rule whose perceived signal–outcome covariances are correct at the present and at the horizon in question, for any law of motion. Series are demeaned throughout, so a constant added to a forecast enters no covariance.

Table 1, Blocks 1 and 2. Each row reports a conditional expectation on some information set, $\mathbb{F}_{i,t} z_{t+h} = \mathbb{E}[z_{t+h} \mid \mathcal{I}_{i,t}]$, so Proposition 1 applies with no further argument. Under full-information rational expectations $\mathcal{I}_{i,t}$ contains z_t. Under sticky or noisy in-

formation (Mankiw and Reis 2002; Woodford 2003; Coibion and Gorodnichenko 2015) it contains a stale or a noisy reading of the state. Under rational inattention (Kohlhas and Walther 2021) the attention paid to each signal is chosen, and in their asymmetric form depends on the state, but the forecaster still conditions on everything she has observed, and Proposition 1 requires only that. Under imprecise memory (da Silveira, Sung and Woodford 2024; Sung 2026) the conditioning set consists of degraded recollections of past signals rather than the signals themselves, and the forecast is the conditional expectation on those recollections. The nonzero entries in the β_{CG} and β_{BGMS} columns are taken from the cited papers.

Table 1, Block 3. The three rows share one structure: the forecaster's estimate of the current state is distorted, and the distorted estimate is carried forward by the true transition. In their common AR(1) environment, therefore, $\mathbb{F}_{i,t} z_{t+h} = \rho^h \mathbb{F}_{i,t} z_t$ up to a constant, and Example 2 gives $\tau_h = 0$ with no linearity assumption; under Assumption 1 the zero extends to any law of motion by Part (ii) of Proposition 2, and Example 4 exhibits the perceived variance $\widetilde{\boldsymbol{\Sigma}}_i$ that represents the distortion as a misperceived precision. What remains is to verify the structure model by model. *Diagnostic Kalman filter* (Bordalo et al. 2020): equation (9) in the text—the nowcast overweights the period-t news by θ, and the h-step forecast is ρ^h times that nowcast. *Ambiguity aversion* (Huo, Pedroni and Pei 2024): their Proposition 1 shows that the ambiguity-averse forecaster's optimal linear strategy equals the Bayesian expectation under a more diffuse prior—an inflated prior variance, which raises the weight on signals—plus a constant pessimistic bias, and in the dynamic case the forecast is the standard Kalman filter applied to that modified problem. The inflated prior variance changes the gain, and so the nowcast; the constant enters no covariance; the transition that carries the nowcast forward is the true one. *Overconfidence in private information* (Adam, Kuang and Xie 2025): the forecaster scales the perceived standard deviation of her private-signal noise down by a constant factor below one and holds rational beliefs about every other parameter, the persistence included (their equation 3.9); her subjective Kalman gain on the private signal is too high, and her h-step forecast iterates her state estimate by ρ^h (their equation C.4).

Table 2, Block 1: fixed rules. The first two rows are Example 2 with a model-specific perceived persistence. *Over-extrapolation* (Angeletos, Huo and Sastry 2020): the forecaster carries the nowcast forward at $\widetilde{\rho} > \rho$, so $\tau_h = \rho^h - \widetilde{\rho}^h < 0$ at every horizon; under-extrapolation, $\widetilde{\rho} < \rho$, gives the opposite sign, and the table reports the direction they estimate. *Cognitive discounting* (Gabaix 2020): the forecaster shrinks the rational forecast of every future deviation from steady state by κ^h with $\kappa \in (0,1)$—his $\bar{m}$—and leaves the nowcast undiscounted; in the AR(1) environment $\mathbb{F}_t z_{t+h} = (\kappa\rho)^h \mathbb{F}_t z_t$, so $\widetilde{\rho} = \kappa\rho$ and $\tau_h = \rho^h(1-\kappa^h) > 0$.

Diagnostic expectations with an undistorted nowcast (Bordalo, Gennaioli and Shleifer 2018): let $z_{t+1} = \rho z_t + \epsilon_{t+1}$ with $\rho \in (0,1)$ and ϵ_{t+1} orthogonal to date-t information, let the nowcast $x_t \equiv \mathbb{F}_t z_t$ be any stationary function of date-t information with $\operatorname{cov}(x_t, z_t) > 0$, and let the forecast tilt the rational forecast from that nowcast by $\theta > 0$ times its revision since $t-1$,

$$\mathbb{F}_t z_{t+h} = \rho^h x_t + \theta\left(\rho^h x_t - \rho^{h+1} x_{t-1}\right). \tag{A2}$$

The actual leg is $\operatorname{cov}(z_{t+h}, x_t) = \rho^h \operatorname{cov}(z_t, x_t)$, because the innovations after t are orthogonal to x_t. In the perceived leg, $\operatorname{cov}(x_{t-1}, z_t) = \rho \operatorname{cov}(x_{t-1}, z_{t-1}) = \rho \operatorname{cov}(x_t, z_t)$, by the orthogonality of ϵ_t to date-$(t-1)$ information and by stationarity, so $\operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t) = \rho^h\left[1 + \theta(1-\rho^2)\right]\operatorname{cov}(x_t, z_t)$. Dividing the difference of the two legs by their common denominator $\operatorname{cov}(x_t, z_t)$,

$$\tau_h = -\theta\,\rho^h\,(1-\rho^2) < 0 \quad \text{at every horizon.} \tag{A3}$$

The nowcast's precision cancels—the wedge is the same under full information, under a rational noisy nowcast, and under a distorted one. What opens it is that the tilt reaches the forecast and not the nowcast; the diagnostic Kalman filter of Example 4 carries the same θ on both and leaves the wedge at zero.

Natural expectations (Fuster, Laibson and Mendel 2010): the forecaster observes the present, fits an AR(1) to z in population, and iterates it, $\mathbb{F}_t z_{t+h} = r_1^h z_t$, where $r_h \equiv$

$\operatorname{cov}(z_{t+h}, z_t)/\operatorname{var}(z_t)$ is the true autocorrelation at lag h. With $\mathbb{F}_t z_t = z_t$, $a_h = r_h$ and $\widetilde{a}_h = r_1^h$, so $\tau_1 = 0$ for any law of motion—the fitted coefficient is the projection of z_{t+1} on z_t—and beyond one step $\tau_h = r_h - r_1^h$ takes the sign of the gap between the true autocorrelation and its geometric extrapolation from the first lag. Under an AR(2) with one-step coefficients (ϕ_1, ϕ_2), the Yule–Walker equations give $r_1 = \phi_1/(1-\phi_2)$ and $r_2 = \phi_1 r_1 + \phi_2$, hence $\tau_2 = \phi_2(1 - r_1^2)$, negative under the momentum-then-reversion dynamics the model is built around. When z is the sum of two orthogonal AR(1) components with persistences $\rho > \delta$ and variance shares ϕ and $1-\phi$, $r_h = \phi\rho^h + (1-\phi)\delta^h$ and $\tau_2 = \phi(1-\phi)(\rho-\delta)^2 > 0$: the fit averages the two persistences and carries the persistent component too little.

Simple models (Molavi 2025): a forecaster constrained to a d-state model tracks the d most persistent components of the economy and treats the rest as i.i.d.

That is his Theorem 5, which holds when $d = 1$ or the lag-one autocovariance matrix is symmetric: the forecast carries each tracked component forward at its own eigenvalue of the lag-one autocorrelation matrix and treats the untracked ones as transitory. The theorem fixes the forecast; the sign of the wedge needs the actual propagation as well, and that depends on the order of the economy's dynamics. When the economy is a first-order vector autoregression whose lag-one autocovariance is symmetric, the components are mutually uncorrelated at every lag and each is an AR(1) at its eigenvalue, so a tracked component is forecast exactly and an untracked one, whose present is observed, has $\mathbb{F}_t z_{t+h} = 0$ for every $h \geq 1$. A surveyed series that loads only on tracked components therefore has $\widetilde{a}_h = a_h$ and $\tau_h = 0$; one that loads on untracked components carries their share of its autocorrelation, $\tau_h > 0$ when their persistence is positive. With higher-order dynamics a tracked component's own autocorrelation departs from the geometric decay at its eigenvalue, and it carries the gap of natural expectations, of either sign. The row reports the first-order case, the multivariate counterpart of the AR(1) environment the other rows use: a zero for a series the model tracks and a positive wedge for one it does not, and CPI inflation's negative wedge fits neither.

Table 2, Block 2: learning. Both rows fall outside Assumption 1, because the loadings respond to realized outcomes, and the wedge is read by the state of learning. *Constant-gain learning* (Evans and Honkapohja 2001): Appendix A3 proves that, when the belief update takes the signal variance as known, the one-step wedge is exactly proportional to the average drift of the loadings—zero once the drift has stopped, for any law of motion and any gain, and proportional to the remaining drift while learning is under way. With the variance learned as well, and beyond one step in either case, the wedge of a settled learner is near zero rather than zero, with an error that shrinks as the gain gets smaller, and the table enters it as ≈ 0. *Bayesian learning* (Farmer, Nakamura and Steinsson 2024): the forecaster models the series as the sum of a random-walk component and an AR(1) component (their equations 15–16), does not know the states or the parameters—the transitory persistence, the random-walk share of the innovation variance, and the innovation scale—and forecasts by Bayes's rule from her current posterior. Her perceived propagation is the propagation implied by that posterior, and until the posterior has converged it is set by her prior: initial beliefs that put too much weight on the random-walk component carry the present forward too far, too little weight too near, and their simulations with upward- and downward-biased initial beliefs produce forecast anomalies of either sign. Because the loadings move with the posterior, the wedge is date-specific here (Section 1.4); the table reports the sign it takes before convergence.

A3. Constant-gain learning

In the second block of Table 2, forecasters fall outside Assumption 1 itself, not just outside the null. A forecaster whose loadings keep responding to realized outcomes needs a result of her own. The leading case is constant-gain least-squares learning, the workhorse of the learning literature (Evans and Honkapohja 2001; Eusepi and Preston 2011). The forecaster runs the regression of the outcome on her observations each period, weighting recent observations more heavily, and forecasts with the coefficients she currently holds.

Proposition A1 (Constant-gain learning)**.** *Suppose the current outcome is among the signals* $\mathbf{S}_{i,t}$*, so that* $\mathbb{F}_{i,t} z_t = z_t$*, and the one-step-ahead loadings are re-estimated from the realized out-*

comes by constant-gain least squares, with the signal variance $\mathbf{\Sigma}_i$ in place of its recursive estimate:

$$\mathbb{F}_{i,t} z_{t+1} = \mathbf{b}_t^\top \mathbf{S}_{i,t}, \qquad \mathbf{b}_{t+1} = \mathbf{b}_t + \gamma\, \mathbf{\Sigma}_i^{-1} \mathbf{S}_{i,t} \big(z_{t+1} - \mathbf{b}_t^\top \mathbf{S}_{i,t} \big), \tag{A4}$$

with gain $\gamma > 0$. Then

$$\tau_1 = \frac{\mathbf{C}_{i,0}^\top \mathbb{E}[\, \mathbf{b}_{t+1} - \mathbf{b}_t \,]}{\gamma \, \operatorname{var}(z_t)}. \tag{A5}$$

Proof. Write $e_{t+1} \equiv z_{t+1} - \mathbf{b}_t^\top \mathbf{S}_{i,t}$ as the forecast error. With $\mathbb{F}_{i,t} z_t = z_t$ and demeaned series, $\tau_1 = \mathbb{E}[z_t e_{t+1}] / \operatorname{var}(z_t)$, and taking expectations in equation (A4) gives $\mathbb{E}[\mathbf{S}_{i,t} e_{t+1}] = \gamma^{-1} \mathbf{\Sigma}_i \mathbb{E}[\mathbf{b}_{t+1} - \mathbf{b}_t]$. Since the present is among the signals, $z_t = \boldsymbol{\beta}_{i,0}^\top \mathbf{S}_{i,t}$ with $\boldsymbol{\beta}_{i,0}^\top \mathbf{\Sigma}_i = \mathbf{C}_{i,0}^\top$; premultiplying the second display by $\boldsymbol{\beta}_{i,0}^\top$ and dividing by $\operatorname{var}(z_t)$ gives equation (A5). □

Equation (A5) is an identity, assuming neither stationarity of her beliefs nor their convergence: it follows from the update rule alone. Two readings follow. A learner whose beliefs have reached their stationary distribution—or merely stopped drifting on average—has learned the propagation of her own signals, whatever levels she settled on: $\tau_1 = 0$ exactly, for any law of motion and any gain. Missing variables do not disturb the zero: settling makes her error orthogonal to every signal she does use, and the present is one of them; omission costs her only beyond one step. Neither does an intercept: constants drop out of every covariance, and a learner who estimates one settles with zero mean error as well. Settled beliefs, then, cannot be the source of a one-step wedge. A learner still moving is reported rather than excused, since τ_1 then measures the present-outcome component of the gap she has left to close: an inability to learn the propagation is a source of a wedge, not a defense against one (Appendix C5). The identity binds at one step.[16] A formal treatment of time-varying perception is left to future work.

[16] With the second-moment normalization itself learned, as in the literature's recursive form, the one-step zero holds to order γ rather than exactly; beyond $h = 1$ a convexity term of order γ enters regardless, with a sign that depends on whether the forecast uses the contemporaneous or the lagged loadings. Only the least-squares form of the update matters: any non-random gain sequence and normalization, decreasing gains included, delivers the same identity date by date. With an intercept the identity holds for the slope block of $\mathbf{b}_t$. Nor need the nowcast be exact: when she learns it from the same signals, the two wedge covariances are equal at her settled point by the symmetry of $\mathbf{\Sigma}_i$—the same mechanism as $\widetilde{\mathbf{\Sigma}}_i$ in Proposition 2, with beliefs generated by an algorithm instead of posited.

APPENDIX B. THE ESTIMATOR AND THE DATA

B1. The estimator

For each variable the estimator stacks, across horizons $h = 1, \ldots, H$, the sample covariances $\operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t)$, $\operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t)$, and the common $\operatorname{cov}(\mathbb{F}_t z_t, z_t)$ as moment conditions, estimated by a single just-identified GMM. The propagation coefficients $a_h, \widetilde{a}_h$, the wedge $\tau_h = a_h - \widetilde{a}_h$, and the across-horizon averages are smooth functions of these moments, with standard errors by the delta method over a Newey–West HAC covariance of the stacked moments. This is the estimator behind every wedge the paper reports, Figure 1 onward; the data construction is in Appendix B2.

Here x is a series as the Survey of Professional Forecasters reports it, z_t is the analyzed series of Table 3, and $\mathbb{F}_t z_{t+h}$ is the consensus (mean) forecast of z_{t+h} made in survey quarter t, with $\mathbb{F}_t z_t$ the nowcast. In change form the nowcast is the survey's current-quarter value less its previous-quarter value, and the forecast of z_{t+h} is the survey's forecast for quarter $t+h$ less its forecast for quarter $t+h-1$. The headline sample is the $T = 140$ forecast origins $t = 1984\text{Q1}, \ldots, 2018\text{Q4}$, so the realizations run through 2019Q4.

Each of z_t, $\mathbb{F}_t z_t$, z_{t+h} and $\mathbb{F}_t z_{t+h}$ is demeaned by its own sample mean, and three products are formed from the demeaned series: $d_t = \mathbb{F}_t z_t \cdot z_t$, $p_{t,h} = z_{t+h} \cdot \mathbb{F}_t z_t$ and $q_{t,h} = \mathbb{F}_t z_{t+h} \cdot z_t$. With bars denoting sample means over t, $a_h = \bar{p}_h / \bar{d}$, $\widetilde{a}_h = \bar{q}_h / \bar{d}$ and $\tau_h = (\bar{p}_h - \bar{q}_h)/\bar{d}$, the sample counterparts of Definition 1; the four-horizon averages replace $p_{t,h}$ and $q_{t,h}$ by their means over h. The two components of equation (11) use in addition $s_{t,h} = z_{t+h} \cdot z_t$, formed in the same way: $r_h = (\bar{s}_h - \bar{p}_h)/\bar{d}$ and $T_h = (\bar{s}_h - \bar{q}_h)/\bar{d}$, so the identity $T_h = \tau_h + r_h$ holds exactly in the sample. Every statistic thus has the form $\bar{n}/\bar{d}$. Its standard error is $\sqrt{\Omega/T}$, where Ω is the Newey–West long-run variance, with Bartlett weights and four lags, of $u_t = (n_t - \bar{d}^{-1}\bar{n}\, d_t)/\bar{d}$. These equal the delta-method standard errors of the just-identified generalized method of moments system in the same moments.

The estimator is consistent, and in samples of the length used here it leans negative. In the rational consensus calibrated to each headline series in Section 2.4, where the pop-

ulation wedge is zero, the mean of the four-horizon average estimate across samples of 140 origins is -0.45, -0.32 and -0.24 of its sampling standard deviation for the unemployment change, the Treasury-bill change and the inflation level; the larger part comes from demeaning a persistent series within the sample and the rest from the ratio's random denominator, and the lean halves when the sample length quadruples. Across the 29 cells of Table 3, each simulated at its own sample length, the lean is negative in 25 and largest in the persistent level cells, from -0.26 to -0.87 standard deviations, where under correct perception a negative estimate appears in 74 to 82 percent of samples for the four persistent rate levels; a count of negative level estimates therefore overstates its evidence, and the paper reports such counts without a probability. Because the lean is negative, it cannot produce a positive wedge: against the calibrated null the one-sided probability of the measured unemployment and Treasury-bill wedges is 0.010 and 0.001, and that of the inflation wedge about 0.07. The Newey–West standard error of the four-horizon average understates its sampling standard deviation by 14 to 19 percent in the headline cells, so those confidence intervals are somewhat narrow.

B2. Data construction

This appendix records how every series of Table 3 is built, so that the headline estimates of Figure 1, the 29-cell panel, and the individual, Greenbook and Livingston estimates of Figure 2 measure the same object and differ only in the forecasts used. Every choice below is taken from the scripts of the replication package, named in Appendix B2.5.

B2.1 The estimand

For each cell the reported average wedge is

$$\overline{\tau} = \frac{1}{H}\sum_{h=1}^{H}\left(a_h - \widetilde{a}_h\right), \qquad a_h = \frac{\operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t)}{\operatorname{cov}(\mathbb{F}_t z_t, z_t)}, \quad \widetilde{a}_h = \frac{\operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t)}{\operatorname{cov}(\mathbb{F}_t z_t, z_t)}, \tag{B1}$$

with every covariance replaced by the sample mean of the product of demeaned series, and its standard error is the Newey–West long-run variance of the ratio's influence function with Bartlett weights and four lags, as Appendix B1 states; these are the delta-method

standard errors of the just-identified generalized method of moments in the same moments. The common denominator $\operatorname{cov}(\mathbb{F}_t z_t, z_t)$ pairs the nowcast with the realized current value, so the two must be built on the same one-quarter footing in every survey.

B2.2 Units: the convention of Table 3

Every series enters in the units of Table 3, on the realized side and on the forecast side alike, and the estimator works in annualized percent, 400 times the quarterly log rates; a common rescaling leaves the wedge, a ratio of covariances, unchanged. Three cases cover the panel.

Inflation rates. The SPF annualizes a quarterly inflation rate by discrete compounding, $g = 100[(P_t/P_{t-1})^4 - 1]$ (survey documentation, pp. 28–29), so the exact inverse that puts a reported rate x in log units is $100\ln(1 + x/100)$, which equals $400\,\Delta \ln P$ for the quarter; it is applied to the nowcast and to every horizon, and the realized side is $400\,\Delta \ln P$ on the second-release vintage. In change form the nowcast is the survey's current-quarter rate less its previous-quarter rate, both mapped, and each forecast is the difference between adjacent horizons. This holds for CPI, core CPI, PCE and core PCE; the GDP deflator, which the survey reports as an index, enters as $400\,\Delta \ln$ of the level on both sides and has no change form.

Interest rates. A rate x is quoted as a simple annual rate, so a quarter's log return at that rate is $\ln(1 + x/400)$; the realized rate and every forecast enter as $400\ln(1 + x/400)$. This is not the inflation map: the survey annualizes inflation by compounding and quotes a rate by simple interest, and each series gets the inverse of its own convention. In change form the nowcast is the survey's current-quarter rate less its previous-quarter rate; for the two corporate yields, whose previous-quarter column is empty in every survey, the lag is taken from the realized yield, public before the survey and never revised—on the 10-year note, where the survey does publish the column, the substitution moves no statistic by more than 0.008. The realized 3-month bill is the secondary-market rate the survey asks for, FRED `TB3MS`, averaged over the quarter; the 10-year note and Moody's Aaa and Baa yields are FRED quarterly averages; none is revised.

The unemployment rate and the quantity variables. The unemployment rate, a share rather

than a return, enters as reported, in level and in change form. Real and nominal GDP, consumption, the two investment aggregates, the two government aggregates, industrial production, housing starts and payroll employment enter as $400\,\Delta \ln$ of a vintage-consistent level on both sides; they have no level form.

B2.3 Realized outcomes: real-time second releases

Realized outcomes are second releases, the value in the second vintage that contains the complete quarter, the convention of Coibion and Gorodnichenko (2015). For the unemployment rate and the price indices the vintages are monthly, from the Real-Time Data Set for Macroeconomists, averaged within the quarter, with inflation computed within that vintage; core CPI has no real-time vintage and uses the latest FRED series. For the trending series the growth rate is the log difference taken *within* the second vintage that contains quarter t, $z_t = 400[\ln L_{t,j} - \ln L_{t-1,j}]$ with j that vintage—quarterly vintages for real and nominal GDP, the deflator and the NIPA components, monthly vintages averaged within the quarter for industrial production, housing starts and payrolls—so that a rebasing between vintages cannot appear as a growth spike, and so that the average of four annualized quarterly log growth rates equals the annualized four-quarter rate exactly, which the Livingston harmonization below relies on. The interest rates are FRED series and are never revised. The four-quarter-ahead realizations of the last origin, 2018Q4, run through 2019Q4.

B2.4 Matching the Livingston Survey to the SPF

The Livingston Survey is semi-annual (June and December) and reports the level of each indicator at the survey month (ZM), six months ahead (6M) and twelve months ahead (12M). Its inputs are built to coincide, term by term, with the quarterly SPF construction, on the survey's own definitions (Philadelphia Fed Livingston documentation, Section 8 and its data descriptions). Table A1 lays out the mapping; four points carry it.

1. *One-quarter nowcast.* The nowcast is the perceived current-quarter move: the forecaster's level at the survey month (or its log, for the CPI) against the realized level three months earlier, and the matching realized current value z_t spans that same quar-

ter. This is the analogue of the SPF current-quarter nowcast.

2. *What the survey asks for.* The 3-month bill is asked for on the secondary-market discount basis, the SPF's basis, but "on the last day of the period," so its realized counterpart is FRED `DTB3` on the last business day of the month rather than a monthly or quarterly average; both the realized rate and the group-mean forecast levels are put on the quarterly-log-return map $400\ln(1 + x/400)$ before any change is formed. The CPI is asked for as a level, not seasonally adjusted in surveys before December 2004 and seasonally adjusted from that survey on; since the realized series and the SPF are seasonally adjusted, each earlier forecast level is multiplied by the seasonal factor of the month it is for, `CPIAUCSL/CPIAUCNS` (FRED, latest vintage), and inflation is then the log change of the level. The unemployment rate is seasonally adjusted and monthly in survey and realization alike.
3. *Real-time vintages.* The realized outcomes, current and future, are the monthly second release from the Real-Time Data Set; the nowcast base, the level three months before the survey that the forecaster differences against, is the first release, the value available at the survey.
4. *Semester leads aggregate to the SPF year.* Each semester lead and forecast is expressed as the average of the two quarterly values it spans: for the CPI, $200\,\Delta\ln$ over the semester, one-half the sum of two annualized quarterly inflations by log-additivity; for the two rates, one-half the six-month change. By linearity of covariances the two-semester average wedge then equals the four-quarter SPF average exactly.

Forecaster-affiliation groups merge the Philadelphia Fed codes: finance (commercial and investment banking), nonfinancial business, and research (academic institutions, consulting firms and government); the remaining small categories enter only the all-forecaster estimate. Each group's forecast is the mean over its members of the reported levels, then differenced or log-differenced.

The rate variables: exact timing, less precision. The match is exact for inflation, a flow. Unemployment and the bill are stocks, for which the two surveys ask for different objects: the SPF the quarterly average, the Livingston Survey the level at a month (the survey

Table A1. Harmonization of the SPF headline and the Livingston panel

	Survey of Professional Forecasters (headline)	Livingston Survey
Frequency	quarterly	semi-annual (June/December)
Units	Table 3: Δx unemployment, $\Delta \ln(1+x\%/4)$ bill, $\ln(1+x\%)/4$ CPI	the same, formed from the reported level fan
Nowcast $\mathbb{F}_t z_t$	current-quarter forecast vs. previous-quarter value	level at the survey month vs. realized level three months earlier
Current z_t	realized current-quarter change / inflation	realized one-quarter change / inflation (same window)
Horizon leads	4 quarters ($h = 1$–4)	2 semesters ($h = 1, 2$), each the average of its two quarters
Reported wedge	$\frac{1}{4}\sum_{h=1}^{4} \tau_h$	$\frac{1}{2}\sum_{h=1}^{2} \tau_h$ (= the SPF average, exactly)
Realization	second release; bill: FRED `TB3MS`, quarterly average	second release; bill: FRED `DTB3`, last business day of the month
Seasonal adjustment	SPF forecasts and realizations seasonally adjusted	CPI forecast levels not seasonally adjusted before the December 2004 survey, converted with `CPIAUCSL/CPIAUCNS`
Population	SPF professionals	Livingston professionals, overall and by affiliation

Notes: The table sets out the construction behind the SPF estimates of Figure 1 and the Livingston Survey estimates of Figure 2. Both use a one-quarter nowcast against real-time vintages and report the average wedge over the year ahead. In the common log units of Appendix B2.2, and with each semester lead the average of its two quarters, the Livingston two-semester average and the SPF four-quarter average estimate the same statistic, so the two differ only in the forecaster population and the sample window.
Sources: as listed in Table A2.

month, +6, +12; for the bill, the last day). Each forecast is set against its realized counterpart at its own timing, so within each survey the realization is timed to the forecast exactly, and using a point-in-time realization for the SPF would introduce a mismatch rather than remove one. What differs across the surveys is precision: with 54 rounds the Livingston rate wedges carry standard errors of 0.07 and 0.10, against 0.04 for inflation. On the common 1992–2018 window the Livingston estimates are 0.17, 0.31 and -0.11 against SPF estimates of 0.22, 0.23 and -0.12, each gap within one Livingston standard error; Section 2.2 reports every Livingston estimate against the SPF interval.

The Greenbook staff panel. The Federal Reserve estimates of Figure 2 use one Greenbook (now Tealbook) projection per quarter: the one prepared for the FOMC meeting whose publication date is nearest the middle of the quarter, the SPF's submission deadline. Its nowcast and one- to four-quarter-ahead projections replace the SPF consen-

sus; they enter in the Greenbook's own units, are mapped exactly as the SPF's (Appendix B2.2), and are set against the same second-release realizations over the same forecast origins, 1984Q1–2018Q4. The staff's 3-month bill projection is published through 2014Q4, so that cell ends there; its bill is, like the SPF's, a quarterly average on the secondary-market discount basis, and its CPI inflation, like the SPF's, a compounded annual rate. Source: the Philadelphia Fed Tealbook (formerly Greenbook) Data Set, released with a five-year lag.

The individual SPF responses. The individual estimates of Figure 2 pool the SPF responses that report the nowcast and all four forecasts for the series, every series demeaned within forecaster, so that a respondent's constant bias enters no covariance and entry and exit cannot contaminate a pooled moment; the standard error is Driscoll–Kraay, the long-run variance over forecast origins of the moments summed within origin, with Bartlett weights and four lags, and the difference from the consensus wedge is computed jointly by stacking the two estimators.

B2.5 Data sources and reproduction

All inputs are public, from the Federal Reserve Bank of Philadelphia and FRED (Table A2).

The build runs in four data-preparation scripts, one estimation script and the exhibit scripts of the replication package. *Realizations:* `build_realtime.py` assembles the second-release realized series from the Real-Time Data Set and FRED, and `build_loggrowth.py` the vintage-consistent log growth of the trending series. *Greenbook:* `build_greenbook.py` selects one projection per quarter and writes the nowcast and four horizons for each variable. *Livingston:* `export_livingston_unified.py`, on the builders of `export_assembled.py`, writes the semester inputs for every variable and affiliation group. *Estimation:* `levels_matrix_all.do` builds every cell of Table 3 on its convention, writes the cell data (`figures/cells/`) and every estimate with its standard error (`levels_matrix_all_b.csv`); `series_table.m` writes Table 3 from the same specification. *Exhibits:* `wedge_headline_unified.m` (Figure 1), `inforig_unified.m` (Figure 3 and the R^2 table), `falsealarm_unified.m` (the calibrated

rational consensus of Section 2.4) with `falsealarm_population_values.m` (its population values), `individual_wedge_unified.m` and `crossforecast_unified.m` (Figure 2 and Figure A2), `projection_bias_matrix_unified.m` (Figure 4), `cg_decontam_unified.m` (Figure 5), `rolling_lambda_unified.m` and `rolling_wedge_unified.m` (Figure 7 and Appendix E3), `breadth_crosssection_unified.m` (Figure 6), `naive_panel_table.m` (Table A5), `falsealarm_bias_tail_check.m` and `falsealarm_percell_bias_check.m` (the estimator's lean, Appendix B1), and `wedge_bootstrap.py` with `matrix_signs_bootstrap.py` for every block-bootstrap probability.

A replicator must reproduce five non-obvious data choices, each detailed above:

- *Units* (Table 3, Appendix B2.2): interest rates as quarterly log returns, $400\ln(1+x/400)$ on both sides; compounded inflation rates as $100\ln(1+x/100)$, the exact inverse of the survey's compounding, against $400\,\Delta\ln P$ realized; the unemployment rate as reported; the quantity variables and the GDP deflator as $400\,\Delta\ln$ of vintage-consistent levels on both sides.
- *Real-time vintages* (Appendix B2.3): realized outcomes are the second release, log growth is differenced within one vintage, the Livingston nowcast base is the first release, and the interest rates come from FRED unrevised.
- *The Livingston definitions* (Table A1): the bill on the last business day of the month (`DTB3`), the CPI level seasonally adjusted before the December 2004 survey, a one-quarter nowcast, and semester leads as the average of their two quarters, so that the two-semester average equals the four-quarter SPF average exactly.
- *Consensus and groups*: the SPF consensus is the survey's mean forecast; each Livingston group forecast is the cross-forecaster mean of reported levels, then differenced or log-differenced, with the affiliation groups finance = I+B, nonfinancial business = S, and research = A+C+G; the individual SPF estimates use only responses that report the nowcast and all four forecasts.
- *Samples*: forecast origins 1984Q1–2018Q4 for the SPF and the Greenbook (Greenbook bill through 2014Q4), 1992–2018 for the Livingston Survey (54 rounds), each cell of the panel on the longest run without a missing quarter (Table 3); realizations run one year past the last origin.

Table A2. Data sources

Series	Source	Detail
SPF forecasts	Federal Reserve Bank of Philadelphia, Survey of Professional Forecasters	mean responses, current quarter and four ahead, for the twenty series of Table 3; individual responses for the three headline series
Greenbook/Tealbook forecasts	Federal Reserve Bank of Philadelphia, Tealbook (formerly Greenbook) Data Set	one projection per quarter, nearest the middle of the quarter; current quarter and four ahead; five-year release lag; bill through 2014Q4
Real-time realizations	Federal Reserve Bank of Philadelphia, Real-Time Data Set for Macroeconomists	monthly vintages of the unemployment rate, the CPI, PCE and core PCE indices, industrial production, housing starts and payrolls; quarterly vintages of GDP, its deflator and the NIPA components; second releases for outcomes, first release for the Livingston nowcast base
Interest rates (realized)	FRED	3-month bill, secondary market: `TB3MS` quarterly average (SPF, Greenbook) and `DTB3` last business day (Livingston); 10-year note `GS10`; Moody's Aaa and Baa; none revised
Core CPI (realized)	FRED, `CPILFESL`	latest vintage; no real-time vintage exists
CPI seasonal factor	FRED, `CPIAUCSL/CPIAUCNS`	converts Livingston CPI forecast levels before the December 2004 survey
Livingston Survey	Federal Reserve Bank of Philadelphia, Livingston Survey	consensus level fan (survey month, +6, +12) and the forecaster-level panel with affiliation codes

APPENDIX C. THE PROPAGATION WEDGE: THEORY

This appendix develops the theory behind Section 1. It records a sufficient condition for the null, works the second-order and two-component examples, proves the measurement-error immunity (Proposition C1), records the cognitive-noise case (Proposition C2), sets down what a nonzero wedge does and does not establish about rationality, and notes aggregation invariance.

C1. A sufficient condition for the null

Propositions 1 and 2 are proved in Section 1.2, and the proofs are not repeated here. This subsection records instead a condition that is simple to check in applications and that each of the canonical cases meets.

The hypothesis of Proposition 2(ii) restricts only how the nowcast is carried forward, not how the present is perceived; a leading way to meet it is to propagate the nowcast by the actual law of motion. If $z_{t+h} = \mathbf{\Phi}^h z_t + \xi_{t+h}$ with ξ_{t+h} orthogonal to date-t information and forecasts iterate the nowcast by the true $\mathbf{\Phi}$, then $\operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t) = \mathbf{\Phi}^h \operatorname{cov}(\mathbb{F}_t z_t, z_t) = \operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t)$, so $\tau_h = 0$ on any information set, and in the multivariate case the cancellation holds coordinate-wise, provided the nowcast–outcome cross-covariance matrix is symmetric—automatic for a scalar surveyed variable. This is the route taken by Example 2, the diagnostic Kalman filter (Example 4), and the cognitive-noise class (Proposition C2): distorted beliefs about the precision of the nowcast—overconfidence, diagnostic expectations—change $\mathbb{F}_t z_t$ but not the operator that carries it forward, and so leave $\tau_h = 0$ with no linearity assumption in these canonical environments (under a linear forecasting rule the zero extends to any law of motion, Proposition 2); imprecise memory as in Proposition C2 is rational filtering of degraded signals, so it needs no linearity at all (Proposition 1).

C2. Worked examples: second-order dynamics and two components

Two examples fill in the cases between the AR(1) and full-information examples of Section 1.1: a second state coordinate, and a second persistence component. Both illustrate the null of Proposition 2.

Example C1 (Second-order dynamics)**.** Let z follow an AR(2), so the future loads on two coordinates of the present, $z_{t+h} = \phi_{1,h} z_t + \phi_{2,h} z_{t-1} + \xi_{t+h}$, with ξ_{t+h} orthogonal to time-t information and $(\phi_{1,h}, \phi_{2,h})$ the top row of the h-step companion matrix. A forecaster reads the present through *arbitrary* information friction—noisy or partial signals—forms her best estimates $\mathbb{F}_t z_t$ and $\mathbb{F}_t z_{t-1}$ of both coordinates, and carries them forward by the true dynamics, $\mathbb{F}_t z_{t+h} = \phi_{1,h} \mathbb{F}_t z_t + \phi_{2,h} \mathbb{F}_t z_{t-1}$. Then

$$a_h = \phi_{1,h} + \phi_{2,h} \frac{\operatorname{cov}(z_{t-1}, \mathbb{F}_t z_t)}{\operatorname{cov}(z_t, \mathbb{F}_t z_t)}, \qquad \widetilde{a}_h = \phi_{1,h} + \phi_{2,h} \frac{\operatorname{cov}(\mathbb{F}_t z_{t-1}, z_t)}{\operatorname{cov}(\mathbb{F}_t z_t, z_t)}, \qquad \tau_h = 0. \tag{C1}$$

The wedge vanishes because the two second-coordinate covariances coincide, $\operatorname{cov}(z_{t-1}, \mathbb{F}_t z_t) = \operatorname{cov}(\mathbb{F}_t z_{t-1}, z_t)$: the realized past comoves with today's nowcast exactly as the estimated past comoves with today's outcome—a symmetry every best linear estimate satisfies. Friction corrupts the forecaster's reading of z_{t-1} as well as z_t, yet it nets out, because the wedge pairs each forecast with a realized outcome symmetrically. A model-free shortcut that skips the cross-covariances and regresses directly does not survive,

$$\begin{aligned} \frac{\operatorname{cov}(z_{t+h}, z_t)}{\operatorname{var}(z_t)} &= \phi_{1,h} + \phi_{2,h} \frac{\operatorname{cov}(z_{t-1}, z_t)}{\operatorname{var}(z_t)} \quad \text{(future on present)}, \\ \frac{\operatorname{cov}(\mathbb{F}_t z_{t+h}, \mathbb{F}_t z_t)}{\operatorname{var}(\mathbb{F}_t z_t)} &= \phi_{1,h} + \phi_{2,h} \frac{\operatorname{cov}(\mathbb{F}_t z_{t-1}, \mathbb{F}_t z_t)}{\operatorname{var}(\mathbb{F}_t z_t)} \quad \text{(forecast on nowcast)}. \end{aligned} \tag{C2}$$

The two share $\phi_{1,h}$ but differ only in the lag coefficient: the raw $\operatorname{cov}(z_{t-1}, z_t)/\operatorname{var}(z_t)$ against $\operatorname{cov}(\mathbb{F}_t z_{t-1}, \mathbb{F}_t z_t)/\operatorname{var}(\mathbb{F}_t z_t)$, the latter inflated by the smaller variance of the smoothed nowcast. Friction drives the two apart, so the shortcut reports a spurious gap, $\phi_{2,h}$ times their difference, exactly where $\tau_h = 0$.[17]

Example C2 (Two-component state)**.** Let the present split into a persistent and a transitory component, $z_t = c_t^P + c_t^T$, two orthogonal AR(1)s with persistences $\rho > \delta$ and

[17] The forecast-on-nowcast slope in fact equals a_h—the forecast is the projection of the future, and the nowcast lies in the information set—so the shortcut errs only in its raw benchmark, which the wedge replaces with $\widetilde{a}_h$. Appendix D2 develops this false positive analytically in a persistent-plus-transitory form, and Section 2.4 gives a calibrated case, the unemployment rate, in which the comparison fires and the wedge does not.

stationary variances V_P, V_T; write $\phi \equiv V_P/(V_P + V_T)$ for the persistent component's share of $\text{var}(z_t)$. Forecasters read the present through a single, possibly noisy nowcast, apportion it between the two components in perceived shares $\widetilde{\phi} : 1 - \widetilde{\phi}$ with $\widetilde{\phi} \equiv \widetilde{V}_P/(\widetilde{V}_P + \widetilde{V}_T)$, and carry each part forward at its own perceived persistence, so that $\mathbb{F}_t z_{t+h} = (\widetilde{\phi}\,\widetilde{\rho}^h + (1-\widetilde{\phi})\widetilde{\delta}^h)\mathbb{F}_t z_t$. Then

$$a_h = \phi\,\rho^h + (1-\phi)\delta^h, \qquad \widetilde{a}_h = \widetilde{\phi}\,\widetilde{\rho}^h + (1-\widetilde{\phi})\widetilde{\delta}^h, \qquad \tau_h = a_h - \widetilde{a}_h. \tag{C3}$$

Here a_h is the true autocovariance shape of z—a blend of the two decay rates weighted by their variance shares—and $\widetilde{a}_h$ is that shape redrawn with the forecaster's perceived rates and perceived split. The nowcast's overall attenuation and noise drop out as in the AR(1) case; so does any common over- or under-confidence that rescales $\widetilde{V}_P$ and $\widetilde{V}_T$ together, since only their ratio—the perceived share $\widetilde{\phi}$—enters $\widetilde{a}_h$. The wedge is therefore blind to the overall precision of the nowcast and sees only the *relative* variance shares of the two components: $\tau_h = 0$ at every horizon exactly when $\widetilde{\rho} = \rho$, $\widetilde{\delta} = \delta$, and $\widetilde{\phi} = \phi$. A wrong rate and a wrong split each open a wedge, and no single horizon pins the perceived mixture down—a too-persistent perceived rate can be offset by a too-small perceived share to match a_h at one horizon, but not at all of them. A misperceived split is itself a misperceived signal–outcome covariance, the general statement of which is Proposition 2.

Natural expectations in both examples. A forecaster who observes the present, fits an AR(1) to z in population, and iterates it—the natural expectations of Fuster, Laibson and Mendel (2010)—forecasts $\mathbb{F}_t z_{t+h} = r_1^h z_t$, where $r_h \equiv \text{cov}(z_{t+h}, z_t)/\text{var}(z_t)$ is the true autocorrelation at lag h. With $\mathbb{F}_t z_t = z_t$, $a_h = r_h$ and $\widetilde{a}_h = r_1^h$, so $\tau_1 = 0$ for any law of motion: the fitted coefficient is the projection of z_{t+1} on z_t, and the one-step propagation of her own signal is learned exactly. Beyond one step, $\tau_h = r_h - r_1^h$ takes the sign of the gap between the true autocorrelation and its geometric extrapolation from the first lag. In Example C1 with one-step coefficients $(\phi_{1,1}, \phi_{2,1})$, the Yule–Walker equations give $r_1 = \phi_{1,1}/(1-\phi_{2,1})$ and $r_2 = \phi_{1,1} r_1 + \phi_{2,1}$, so $\tau_2 = \phi_{2,1}\,(1 - r_1^2)$: negative when $\phi_{2,1} < 0$, the momentum-then-reversion dynamics natural expectations are built around, where the

fit carries a shock too far. In Example C2, $r_h = \phi\rho^h + (1-\phi)\delta^h$ with the share ϕ defined there, and $\tau_2 = \phi(1-\phi)(\rho-\delta)^2 > 0$: the fit averages the two persistences and carries the persistent component too little. The tracked component of a simple model (Molavi 2025) is this same AR(1) fit: his Theorem 5 carries each tracked component forward at its own eigenvalue of the lag-one autocorrelation matrix, which is that component's r_1, so a tracked component that is not itself an AR(1) carries the gap $r_h - r_1^h$ of either sign, a tracked AR(1) component carries none, and the variables the model does not track, which it forecasts as i.i.d., carry $\tau_h = a_h$; the simple-models row of Table 2 reports the last two cases, the ones a first-order economy admits.

C3. Measurement-error immunity

Proposition C1 (Classical measurement-error immunity)**.** *If the recorded forecasts and nowcast equal the true ones plus mean-zero noise uncorrelated with the realized series, then every covariance in equation (2) is unchanged, and the wedge is identical in population.*

Proof. Let the recorded forecasts and nowcast be $\widehat{\mathbb{F}}_t z_{t+h} = \mathbb{F}_t z_{t+h} + \eta_{h,t}$ and $\widehat{\mathbb{F}}_t z_t = \mathbb{F}_t z_t + \eta_{0,t}$, with each error mean-zero and uncorrelated with the realized series $\{z_t, z_{t+h}\}$. Each covariance in equation (2) pairs a forecast or realization with the realized state z_t, or—through the recorded nowcast—with z_t and z_{t+h}. By assumption every η is orthogonal to those realized series and so contributes zero: $\operatorname{cov}(z_{t+h}, \widehat{\mathbb{F}}_t z_t) = \operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t)$, $\operatorname{cov}(\widehat{\mathbb{F}}_t z_{t+h}, z_t) = \operatorname{cov}(\mathbb{F}_t z_{t+h}, z_t)$, and the shared denominator $\operatorname{cov}(\widehat{\mathbb{F}}_t z_t, z_t) = \operatorname{cov}(\mathbb{F}_t z_t, z_t)$. The noise in the recorded *nowcast* is differenced out for the same reason: it is orthogonal to both realized legs. □

Remark. The immunity is special to the wedge's construction from forecast-with-realization covariances. The individual-level overreaction regression of Bordalo et al. (2020) regresses forecast errors on forecast *revisions*; the revision is a difference of two forecast objects, so the noise variance enters the slope and biases it toward apparent overreaction, and—unlike the attenuation in the consensus Coibion–Gorodnichenko regression—the bias does not vanish in a panel.

C4. The cognitive-noise class

Proposition C2 (Rational-inattention / cognitive-noise null)**.** *Consider a forecaster who tracks a latent, possibly multi-component state evolving by a true transition* $\mathbf{\Phi}$*, observing it through capacity-constrained and possibly memory-degraded Gaussian signals, and forming the posterior-mean forecast. Then her h-step forecast iterates her nowcast by* $\mathbf{\Phi}^h$*, and the consensus—an average of such forecasts—does likewise; hence* $\tau_h = 0$*.*

Proof sketch. With Gaussian signals and a linear transition the posterior mean is a linear projection of the state onto the information set, and the h-step posterior mean is $\mathbf{\Phi}^h$ times the current posterior mean, whatever the signal precisions or capacity allocation; the per-forecaster hypothesis of Proposition 2 therefore holds. The consensus is a cross-forecaster average of these projections, itself $\mathbf{\Phi}^h$ times the average current posterior mean, so the iteration—and the null—survives aggregation. □

C5. What a nonzero wedge says about rationality

The null of Proposition 2 is a necessary, not a sufficient, condition for rational expectations, and it is worth setting down precisely what its rejection does and does not establish. A forecaster clears the null whenever her perceived signal–outcome covariance is correct, $\widetilde{\mathbf{C}}_{i,0} = \mathbf{C}_{i,0}$ and $\widetilde{\mathbf{C}}_{i,h} = \mathbf{C}_{i,h}$, and rational expectations imply this on any information set, so a nonzero wedge rejects rational expectations without naming that set (Proposition 1). What it rejects is in fact the wider class of all correctly perceived propagation profiles, of which rational expectations is one member: a rejection rules out correct propagation, not optimization, which Proposition 2(i) shows can never be tested.

Limited information is not, on its own, a violation. The identifying argument holds on any information set, so a forecaster who reads the present through few, noisy, or selected signals—yet carries whatever she sees forward by the true dynamics—leaves $\tau_h = 0$. Her ignorance of the present surfaces not in the wedge but in the information-friction component r_h of equation (11), which the decomposition holds apart from τ_h. What a forecaster knows can never generate a wedge (Proposition 2)—so the wedge speaks to whether she propagates what she sees correctly, not to how much she sees.

A rejection establishes one thing: for at least one forecaster, no admissible representation matches the true signal–outcome covariances at horizons 0 and h—a misperceived signal–outcome covariance, the present mapped to the future by the wrong dynamics. In the univariate reading of Example 2 this is $\tau_h = \rho^h - \widetilde{\rho}^h$, the gap between true and perceived persistence; in the two-component state of Example C2 it is a wrong perceived rate or a wrong perceived split, the latter itself a misperceived covariance. Errors that leave perceived propagation correct do not register, however non-rational they may be: a distorted belief about the *precision* of the signal—overconfidence, diagnostic expectations—distorts only the perceived precision $\widetilde{\boldsymbol{\Sigma}}_i$, leaves the perceived signal–outcome covariances correct, and so cancels by the symmetry of Proposition 2(ii) (Example 4), imprecise memory leaves the forecast a conditional expectation on the degraded information and so cancels by the tower property (Propositions 1 and C2), as does a pure level bias, while a merely coarse information set falls into r_h.[18] The wedge sees only the propagation slice of any departure from rationality: $\tau_h = 0$ fails to reject correct propagation but certifies neither it nor rational expectations—offsetting misperceptions can also deliver a zero wedge at a given horizon, as in Example C2—while $\tau_h \neq 0$ identifies misperceived propagation and nothing else.

This speaks to a natural objection: that a forecaster who cannot learn the dynamics, or who conditions on the wrong information without realizing it, ought not to be called non-rational. Under rational expectations in the sense of Muth (1961), the subjective distribution must equal the objective one conditional on the information actually used, so a forecaster whose perceived signal–outcome covariance differs from the truth holds beliefs that do not, and departs from rational expectations whether the error came from limited cognition or from an unrepresentative sample. The dividing line is awareness, not breadth. A forecaster who knows her signal is selected or mismeasured and corrects for it—reweighting for the selection, deconvolving the noise—recovers the correct condi-

[18] In the scalar case with a single signal, $\tau_h \propto \widetilde{C}_{i,0}\,\rho^h - \widetilde{C}_{i,h}$, which vanishes precisely when the perceived propagation $\widetilde{C}_{i,h}/\widetilde{C}_{i,0}$ equals the true ρ^h; the perceived signal variance $\widetilde{\Sigma}_i$ cancels between the wedge's numerator and denominator, so a precision error—overconfidence—is invisible to it, while a pure level bias adds a constant of zero covariance; neither disturbs the ratio, and so neither disturbs the wedge. One qualification: a single even horizon cannot separate $\widetilde{\rho} = -\rho$ from $\widetilde{\rho} = \rho$, since $\rho^h = (-\rho)^h$; reading the wedge across horizons, with at least one odd h, removes the knife-edge.

tional beliefs and clears the null with limited and even biased information; one who treats the same signal as representative does not. Rational expectations permit a forecaster to know little, but not to be confidently wrong about what her data mean. An inability to learn the propagation, or the unaware use of wrong information, is therefore a *source* of a nonzero wedge and not a defense against one: it is a way of misperceiving propagation, which the wedge reports—provided the misperception leaves its mark on propagation rather than on precision alone.

A related objection reads the reports rather than the beliefs: a rational forecaster may report something other than her conditional mean—shading it under asymmetric loss, or smoothing it for reputational reasons. Constant shading is differenced out, since the wedge is built from demeaned covariances. Uniform smoothing is inside the null outright: a report that mixes the current rational forecast with its own lags by fixed weights is a fixed-weight combination of rational vintage expectations, and the tower-property identity behind Proposition 1 applies vintage by vintage—each lagged forecast is a conditional expectation on its own information set—so $\tau_h = 0$ with no linearity. What survives is a reporting rule whose distortion varies with the state or the horizon—shading proportional to state-dependent uncertainty under asymmetric loss, say—a nonlinear, horizon-dependent reporting rule outside the maintained linear class, observationally equivalent to misperceived propagation of the reported series. Two findings discipline that reading: the level–change matrix (Section 3.2), which a loss-based account would need to reproduce cell by cell with one parameter per series, and the Greenbook replication (Section 2.2), where the staff forecast faces no publication or reputational incentive.

C6. Aggregation invariance

The numerator covariances in equation (2) are linear in the forecast operator, so $\mathrm{cov}(\mathbb{F}_t z_{t+h}, z_t) = N_t^{-1} \sum_i \mathrm{cov}(\mathbb{F}_{i,t} z_{t+h}, z_t)$ and likewise for the realized leg; the consensus wedge is thus the same statistic whether built from consensus forecasts or from the individual panel pooled over i. The identity presumes the same respondents behind the nowcast leg and the forecast leg—automatic when each respondent answers both items in the same survey round; with item non-response and heterogeneous forecaster

precision, per-item consensus means can open a survey-construction gap, so a balanced-respondent construction is the conservative check. The relation to the individual wedges is a weighted average, not an identity: writing $\tau_{i,h}$ for the wedge built from forecaster i's own forecasts and $d_i \equiv \operatorname{cov}(\mathbb{F}_{i,t} z_t, z_t)$ for its denominator, a fixed set of respondents gives $\tau_h = \sum_i d_i \tau_{i,h} / \sum_i d_i$, so the consensus wedge need not equal the mean of the individual wedges, and the implication runs one way only: a zero wedge for every forecaster makes the consensus wedge zero, while a zero consensus wedge can hide individual departures of opposite sign. When every $d_i > 0$, the weights are positive and sum to one, so $|\tau_h| \leq \max_i |\tau_{i,h}|$: the consensus wedge is a statement about the consensus—the object the SPF measures—and is conservative for individual misperception.

Appendix D. Further results on the wedge and the anchor weight

This appendix supports Sections 2 and 3. It measures the one channel through which an information friction can reach the wedge's size, shows analytically that the persistence comparison can read an information friction as misperception while the wedge cannot, records what the length of the rolling window buys, and sets the anchor weight against two features of the environment.

D1. How much of the wedge's size the friction can carry

The wedge is zero under correct perception for any information set (Proposition 2), but away from that null nothing forces its *size* to be free of what forecasters observe. This subsection isolates the one channel through which the friction can reach that size, and measures it.

The actual leg of the wedge, $a_h = \operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t) / \operatorname{cov}(\mathbb{F}_t z_t, z_t)$, conditions on the nowcast. Write the same quantity conditioned on the realized present,

$$a_h^\circ \equiv \frac{\operatorname{cov}(z_{t+h},\, z_t)}{\operatorname{var}(z_t)}, \tag{D1}$$

and let $\iota_h \equiv a_h - a_h^\circ$ be the gap between the two. Adding and subtracting a_h° splits the

wedge identically,

$$\tau_h = \left(a_h^\circ - \widetilde{a}_h\right) + \iota_h, \tag{D2}$$

with no assumption beyond the finite second moments of Assumption 1.

Two facts make ι_h the friction's only route into the wedge's size. Under full information the nowcast is z_t, so $a_h = a_h^\circ$ and $\iota_h = 0$, for any law of motion. And in the univariate state of Example 2, $z_{t+h} = \rho^h z_t + \xi_{t+h}$ with ξ_{t+h} orthogonal to date-t information gives $\operatorname{cov}(z_{t+h}, \mathbb{F}_t z_t) = \rho^h \operatorname{cov}(z_t, \mathbb{F}_t z_t)$ for any nowcast, so $a_h = \rho^h = a_h^\circ$ and $\iota_h = 0$ for any information set. A nonzero ι_h therefore needs both an imperfect nowcast and dynamics richer than a univariate state: the friction cannot create the wedge, and it can move the wedge's size only here.

On the balanced 1984Q1–2018Q4 sample the across-horizon average $\bar{\iota}$ is 0.03 (standard error 0.04) for the unemployment change, -0.04 (0.03) for the Treasury-bill change and 0.10 (0.04) for CPI inflation. Removing it leaves $\bar{\tau} - \bar{\iota}$ at 0.21, 0.25 and -0.23, against measured wedges of 0.24, 0.21 and -0.13. The only cell in which $\bar{\iota}$ is distinguishable from zero is inflation, and there the wedge net of the gap is nearly twice as large. Classical measurement error in the realized series attenuates a_h°, whose denominator is $\operatorname{var}(z_t)$, and so inflates $\bar{\iota}$; the second-release vintage limits that, and its direction is conservative for the reading just given.

What equation (D2) does not deliver is a split of the wedge into a misperception part and a friction part. Its first term pairs actual propagation measured against the present with perceived propagation measured against the nowcast, and it is not itself zero under the null—only the sum is. Reading it as the wedge the same forecasts would show under full information would require forecasts to iterate the nowcast at a single perceived persistence, $\widetilde{a}_h = \widetilde{a}_1^h$, which the perceived propagation of Figure 1 does not support: the Treasury bill's turns negative by $h = 3$ where iteration implies $+0.03$, and inflation's is nearly flat at 0.38, 0.34 and 0.35 where iteration implies 0.23, 0.11 and 0.05. The wedge of Definition 1 remains the reported object; ι_h bounds how much of its size the friction can carry.

D2. The persistence comparison beyond an AR(1)

Example 1 shows the persistence comparison of Section 2.4 firing under correct perception when the law of motion is an AR(2) and the nowcast lags the outcome. This subsection gives a second case, in which the dynamics have two components and the nowcast is noisy, and shows why the wedge stays at zero in it. Write $\rho_h^* \equiv \operatorname{cov}(z_{t+h}, z_t)/\operatorname{var}(z_t)$ for the realized autocorrelation and $\hat{b}_h^P \equiv \operatorname{cov}(\mathbb{F}_t z_{t+h}, \mathbb{F}_t z_t)/\operatorname{var}(\mathbb{F}_t z_t)$ for the slope of the forecast on the nowcast, so that the persistence comparison is $\rho_h^* - \hat{b}_h^P$.

Take a two-component law of motion,

$$z_t = c_t + \mu_t, \qquad c_t = \rho\, c_{t-1} + \epsilon_t, \qquad \mu_t = \bar{\rho}\, \mu_{t-1} + \eta_t, \qquad 0 < \rho < \bar{\rho} < 1, \tag{D3}$$

with $\epsilon_t \perp \eta_t$ i.i.d. and persistent-component variance share $\theta \equiv \operatorname{var}(\mu_t)/\operatorname{var}(z_t)$. The realized autocorrelation is the variance-weighted mixture $\rho_h^* = (1-\theta)\rho^h + \theta\bar{\rho}^h$; being a mixture of two geometric terms, it is matched by no single AR(1). To isolate the friction, let the forecaster know the past components $\{c_s, \mu_s\}_{s<t}$ *without error* but observe the current state only through a noisy signal $s_t = z_t + \nu_t$, $\nu_t \sim (0, \sigma_\nu^2)$ i.i.d. Knowing the past, she predicts the current components by ρc_{t-1} and $\bar{\rho}\mu_{t-1}$ and revises them from the signal surprise $\tilde{s}_t \equiv s_t - \rho c_{t-1} - \bar{\rho}\mu_{t-1}$ by the Bayesian gains $g_c = \sigma_\epsilon^2/\Sigma$, $g_\mu = \sigma_\eta^2/\Sigma$ with $\Sigma \equiv \sigma_\epsilon^2 + \sigma_\eta^2 + \sigma_\nu^2$. Write the total gain

$$g \equiv g_c + g_\mu = \frac{\sigma_\epsilon^2 + \sigma_\eta^2}{\sigma_\epsilon^2 + \sigma_\eta^2 + \sigma_\nu^2} \in (0, 1], \tag{D4}$$

$g = 1$ being full information. Perception is correct: she uses the true $\rho, \bar{\rho}, \theta$ and gains, so $\mathbb{F}_t z_t = \hat{c}_t + \hat{\mu}_t$ and $\mathbb{F}_t z_{t+h} = \rho^h \hat{c}_t + \bar{\rho}^h \hat{\mu}_t$ with $\hat{c}_t = \rho c_{t-1} + g_c \tilde{s}_t$ and $\hat{\mu}_t = \bar{\rho}\mu_{t-1} + g_\mu \tilde{s}_t$.

Proposition D1. *In the environment of equations (D3) to (D4), with no misperception of the law of motion, the slope of the forecast on the nowcast is itself a mixture,*

$$\hat{b}_h^P = (1-\theta')\rho^h + \theta'\bar{\rho}^h, \qquad \theta' = \frac{\theta\, \kappa_\mu}{(1-\theta)\, \kappa_c + \theta\, \kappa_\mu}, \tag{D5}$$

with $\kappa_c = \rho^2 + g(1-\rho^2)$ *and* $\kappa_\mu = \bar{\rho}^2 + g(1-\bar{\rho}^2)$. *The persistence comparison is therefore*

nonzero under correct perception,

$$\rho_h^* - \hat{b}_h^P = (\theta - \theta')\,(\bar{\rho}^{\,h} - \rho^h), \qquad \theta' - \theta = \frac{\theta(1-\theta)(\bar{\rho}^2 - \rho^2)(1-g)}{(1-\theta)\kappa_c + \theta\kappa_\mu} \;\geq\; 0, \tag{D6}$$

with equality if and only if $g = 1$ (full information), $\rho = \bar{\rho}$, or $\theta \in \{0, 1\}$ (a single component). The perceived slope is inflated ($\theta' \geq \theta$), so the comparison is negative, and its magnitude increases in the friction $1 - g$.

Proof. The prediction terms $\rho c_{t-1}, \bar{\rho}\mu_{t-1}$ and the surprise $\tilde{s}_t$ are mutually orthogonal, with variances $\rho^2 V_c$, $\bar{\rho}^2 V_\mu$, and Σ, where $V_c = \mathrm{var}(c_t)$, $V_\mu = \mathrm{var}(\mu_t)$, $g_c\Sigma = \sigma_\epsilon^2 = (1 - \rho^2)V_c$, and $g_\mu\Sigma = (1 - \bar{\rho}^2)V_\mu$. Substituting $\hat{c}_t, \hat{\mu}_t$,

$$\mathrm{cov}(\mathbb{F}_t z_{t+h}, \mathbb{F}_t z_t) = \rho^h\left[\rho^2 V_c + g(1-\rho^2)V_c\right] + \bar{\rho}^h\left[\bar{\rho}^2 V_\mu + g(1-\bar{\rho}^2)V_\mu\right] = \rho^h V_c\kappa_c + \bar{\rho}^h V_\mu\kappa_\mu,$$

and $\mathrm{var}(\mathbb{F}_t z_t) = V_c\kappa_c + V_\mu\kappa_\mu$ (the $h = 0$ case). Dividing gives equation (D5) with $\theta' = V_\mu\kappa_\mu / (V_c\kappa_c + V_\mu\kappa_\mu)$; subtracting it from $\rho_h^* = (1-\theta)\rho^h + \theta\bar{\rho}^h$ and using $\kappa_\mu - \kappa_c = (\bar{\rho}^2 - \rho^2)(1-g)$ gives equation (D6). □

Mechanism. The noisy nowcast is a smoothed reading of z_t. Smoothing discounts the transitory component—most of whose current movement is unforecastable noise—more heavily than the persistent one, so the nowcast, and the forecasts iterated from it, over-represent the persistent component: $\theta' > \theta$. A regression of the forecast on this persistence-tilted nowcast returns an inflated persistence, while the realized autocorrelation ρ_h^* carries the true weights. Their gap is the false positive. Neither ingredient alone suffices: with one component there is nothing to re-weight ($\rho = \bar{\rho}$), and with full information the nowcast equals z_t and the weights are untouched ($g = 1$).

The wedge in the same environment. Under correct perception both legs of the wedge collapse to the *same* object as the slope of the forecast on the nowcast,

$$a_h = \widetilde{a}_h = \hat{b}_h^P = (1-\theta')\rho^h + \theta'\bar{\rho}^h,$$

each carrying the identical, friction-distorted weight θ' (the equality $a_h = \widetilde{a}_h$ is Proposi-

tion 2). The friction enters both legs in exactly the same amount, so the wedge differences it out: $\tau_h = a_h - \widetilde{a}_h = 0$. The persistence comparison fails for the opposite reason: it pairs the slope $\hat{b}_h^P$ (weight θ') with the raw autocorrelation ρ_h^* (weight θ), and the two weightings differ by precisely the friction.

D3. The length of the window

The windows of Figure 7 hold 60 forecast origins, three times the five years over which Coibion and Gorodnichenko (2015) measure the volatility in their Figure 3. The two are not the same object. Their coefficient is estimated for each quarter separately, from that quarter's cross-section of survey variables and horizons, and the sequence is then smoothed; their twenty quarters is the moving window of the volatility line they plot against it, not an estimation window. The wedge is a ratio of covariances taken over forecast origins, so it needs origins, and the question is how few.

Table A3 answers it under the constant-weight null of Appendix E3, in which $\tau_h = \lambda g_h$ holds exactly, so every movement a window estimate shows is the estimator. What the window buys is the denominator, not the wedge. Going from 20 origins to 60 narrows the null spread of the wedge itself by about a third in all three cells, and narrows the spread of the implied weight by a factor of five in the two change-form cells. In change form the implied weight divides by the window's own actual propagation $\bar{a}$, which at twenty origins wanders close to zero: on the data the estimate then ranges over $[-3.1, 8.7]$ for the unemployment rate and $[-122, 17]$ for the Treasury bill. In level form the divisor is $1 - \bar{a}$, bounded away from zero, and the level cell is usable at twenty origins. Sixty is therefore close to the shortest window on which the implied weight can be read at all. It is not a statement about the wedge, which the shorter windows estimate about as well.

D4. The anchor weight and the environment

Figure A1 adds to Figure 7 a row for the two candidate features of the environment, and extends every column back as far as its own survey question allows. The volatility of the series within the window is the environment variable of Coibion and Gorodnichenko (2015), measured here on the series being forecast; the mean of the series' own level over

Table A3. What the window length buys

	Wedge $\overline{\tau}$				Implied weight λ			
Origins in the window	20	40	60	80	20	40	60	80
Unemployment rate, change	0.83	0.63	0.53	0.46	9.62	4.98	1.74	1.09
3-month T-bill rate, change	0.69	0.56	0.47	0.41	5.15	3.38	1.22	0.87
CPI inflation rate, level	0.55	0.44	0.37	0.33	0.46	0.39	0.34	0.29

Notes: Each entry is the width of the 5th–95th percentile range of the window estimate under the constant-weight null of Appendix E3, in the units of the statistic; a smaller number is a tighter estimator. The null is the calibrated rational consensus of Section 2.4 with projection bias added at each series' 1984–2018 weight (0.52, 0.77 and 0.17), so that $\tau_h = \lambda g_h$ holds in population and the weight is constant by construction; the table pools all windows of the stated length across 2,000 samples of 140 origins. Both statistics are those of Figure 7: $\overline{\tau}$ is the average wedge over $h = 1, \ldots, 4$, and λ its ratio to the average factor g, $\bar{a}$ in change form and $-(1-\bar{a})$ in level form. Computed by `empirics/rolling_window_length.m`.

the same window is the second candidate, which inflation answers to and the other two do not. Both are measured on the same 60 origins as the weight, so that the environment covers the span the estimate covers.

The correlation between the unemployment weight and the volatility of the series is -0.94 within the estimation sample, -0.63 across all 159 windows and -0.37 across the 142 that do not contain the pandemic. The volatility terciles of the 159 give weights of 0.68, 0.51 and 0.52, so the quiet third stands apart and the other two do not differ; the ratio-free regression of the wedge on the factor it multiplies and on that factor times the standardized volatility gives -0.15 per standard deviation over all the windows and -0.08 over the pre-pandemic ones. The windows reaching before 1984 carry an average weight of 0.64 at a volatility of 0.40, against 0.45 at 0.35 for those ending 2009–2018Q4. For the Treasury bill and CPI the within-sample correlations with volatility are $+0.49$ and -0.28, which the constant-weight null of Appendix E3 produces 34 and 48 percent of the time; the Treasury bill's weight departs from its own average by at most 0.22 over every window the survey supports, against a median single-window standard error of 0.11. Inflation's weight falls with the window mean of inflation across the 23 windows that reach past 2018Q4, correlation -0.94 and ratio-free slope -0.12 per standard deviation of the level (one-sided $p = 0.004$ against the constant-weight null); those windows all contain the 2021–2023 surge, and six series-by-environment combinations were examined, so this is one episode and not twenty-three.

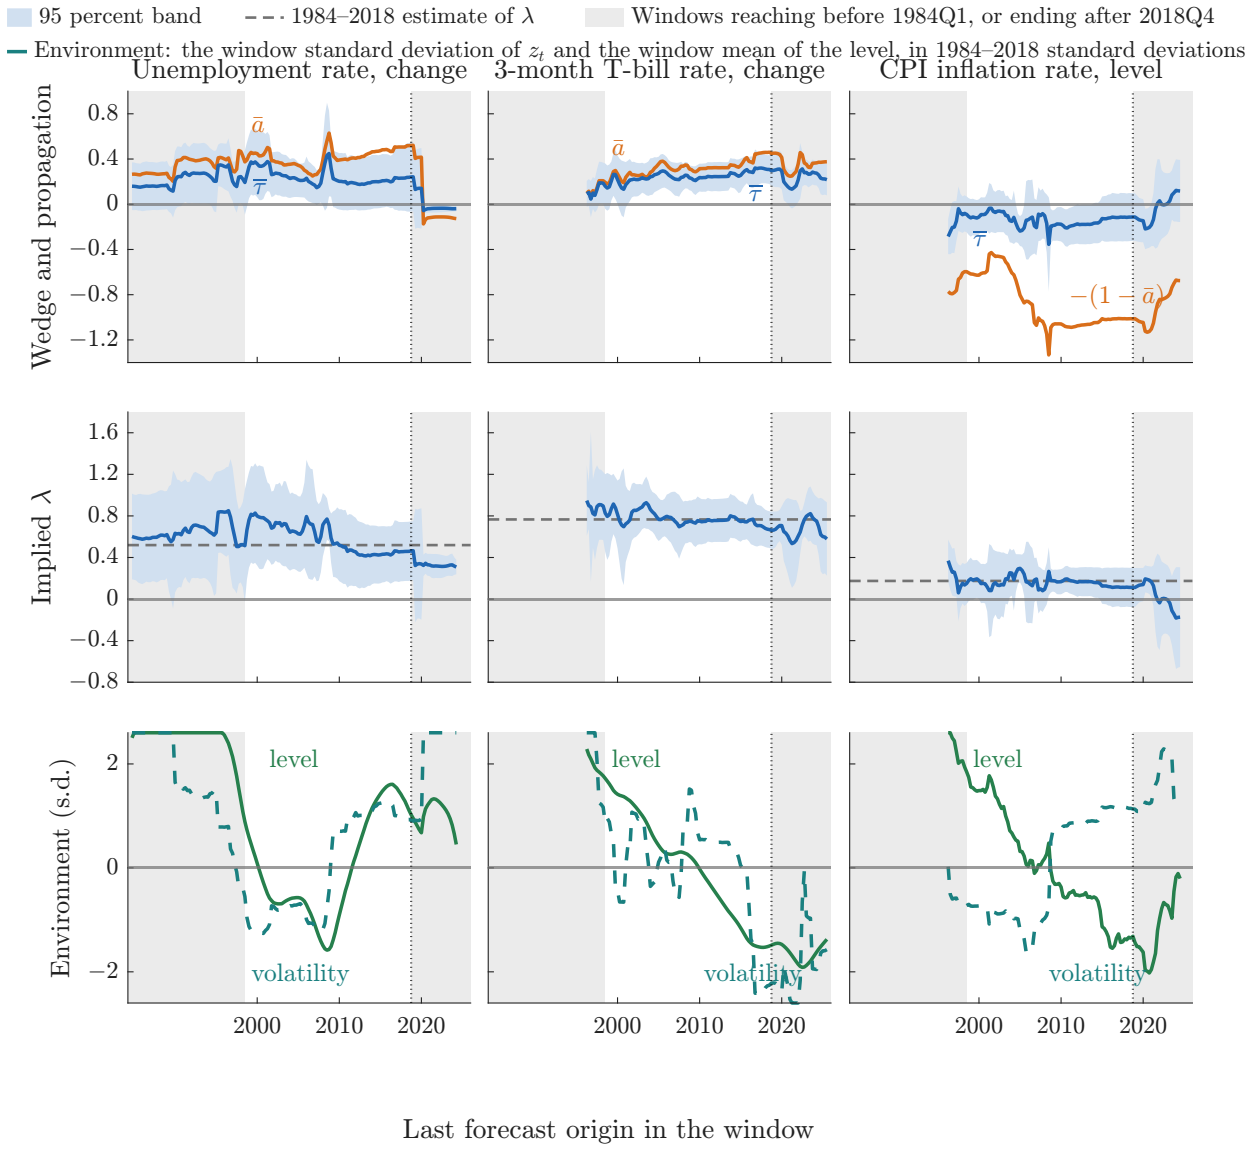


Figure A1. The anchor weight against two features of the environment

Notes: The first two rows are those of Figure 7. The third row plots the standard deviation of z_t within the window (dashed) and the mean of the series' own level over the same window (solid), each in standard deviations of the 81 windows whose last origin runs from 1998Q4 to 2018Q4; for a change-form series the level is the underlying rate, not the mean of z_t, which is a near-zero drift. Every column runs back to the earliest 60-origin window its survey question supports: 1984Q4 for the unemployment rate, asked since 1968Q4, and 1996Q2 for the Treasury bill and CPI, asked since 1981Q3. A window is 60 complete origins; five early unemployment rounds lack the four-quarter-ahead forecast and are treated as missing, as in footnote 5 of Coibion and Gorodnichenko (2015), so two of the 159 windows span one quarter more than the rest. Windows on the gray ground either reach before 1984Q1, the Great Moderation break of McConnell and Perez-Quirós (2000), or end after 2018Q4; the estimation sample is the white region between them. Standardizing on the 81 within-sample windows puts the three columns on one axis but sends two series off it: the unemployment volatility reaches 17 standard deviations in the windows that contain 2020, and the Treasury bill's 8.9 in those that reach 1981–1983. Appendix D4 reports the tests of the relation between the weight and each feature.

APPENDIX E. SUPPLEMENTS TO THE EMPIRICAL SECTIONS

This appendix collects the exhibits that Sections 2 and 3 point to: the coverage check behind the magnitude comparison, the individual SPF responses horizon by horizon, the rolling sample windows, the measured wedge and the three rival diagnostics in every cell of the panel, the level–change test on every series that supports both forms, the test of whether the anchor weight is one number, the anchor's effect on the mean squared forecast error, the robustness of the error–revision check with a replication of the benchmark regression, and the note on why the weight cannot be estimated by regressing forecast errors on forecasts. Every exhibit follows the conventions of Table 3 and the estimator of Figure 1.

E1. Supplement to the magnitude comparison

Table A4 documents the coverage check of Section 2.3: how much of the predictable forecast error the realized present leaves out.

E2. The individual SPF responses, horizon by horizon

Figure A2 reports the comparison of Section 2.2 horizon by horizon: the wedge on the individual SPF responses beside the consensus wedge of Figure 1, and the difference between them, with a standard error computed jointly.

E3. Rolling sample windows

Is the headline pattern of Figure 1 an artifact of the sample window? Re-estimated on every window of 60 consecutive forecast origins within 1984–2018 (the top row of Figure 7 plots them), each series keeps its sign in all 81 windows: the unemployment wedge ranges from 0.15 to 0.45, the Treasury-bill wedge from 0.13 to 0.32 and the inflation wedge from -0.35 to -0.03. Precision is another matter, since a window holds less than half the sample: the wedge is significant in 93 percent of the windows for the Treasury bill but in 46 and 31 percent for unemployment and inflation. Because the windows overlap, this shows stability, not additional evidence. With 40-quarter windows the sign holds in every window for the two rates and in 98 percent of them for inflation. The windows that

Table A4. Regressions of the consensus forecast error on the present and on the revision

	$h=1$	$h=2$	$h=3$	$h=4$
Unemployment rate, change				
R^2: error on outcome	0.10	0.16	0.09	0.11
R^2: error on revision	0.08	0.14	0.05	n/a
R^2: error on both	0.12	0.21	0.11	n/a
3-month Treasury-bill rate, change				
R^2: error on outcome	0.15	0.10	0.09	0.03
R^2: error on revision	0.04	0.01	0.01	n/a
R^2: error on both	0.15	0.11	0.09	n/a
CPI inflation rate, level				
R^2: error on outcome	0.01	0.02	0.01	0.01
R^2: error on revision	0.00	0.00	0.00	n/a
R^2: error on both	0.01	0.02	0.01	n/a

Notes: Series, transformations, sample and sources are those of Figure 1, and $e_{t,h} = z_{t+h} - \mathbb{F}_t z_{t+h}$ is the consensus forecast error at horizon h. Each panel reports the R^2 of ordinary least squares regressions, each with an intercept, of $e_{t,h}$ on the realized current value z_t (error on outcome), on the consensus forecast revision $\mathbb{F}_t z_{t+h} - \mathbb{F}_{t-1} z_{t+h}$, the regressor of Coibion and Gorodnichenko (2015) (error on revision), and on both. The two forecasts in a revision are of the same quarter, $t+h$, made in consecutive surveys; in change form each is the survey's forecast for quarter $t+h$ less its forecast for quarter $t+h-1$. The increase in R^2 from adding the revision measures how much of the predictable error the present leaves out; it does not compare the two channels, because misperceived propagation also makes the revision predict the error. At $h=4$ the revision would require a five-quarter-ahead forecast, which the SPF does not report (n/a).

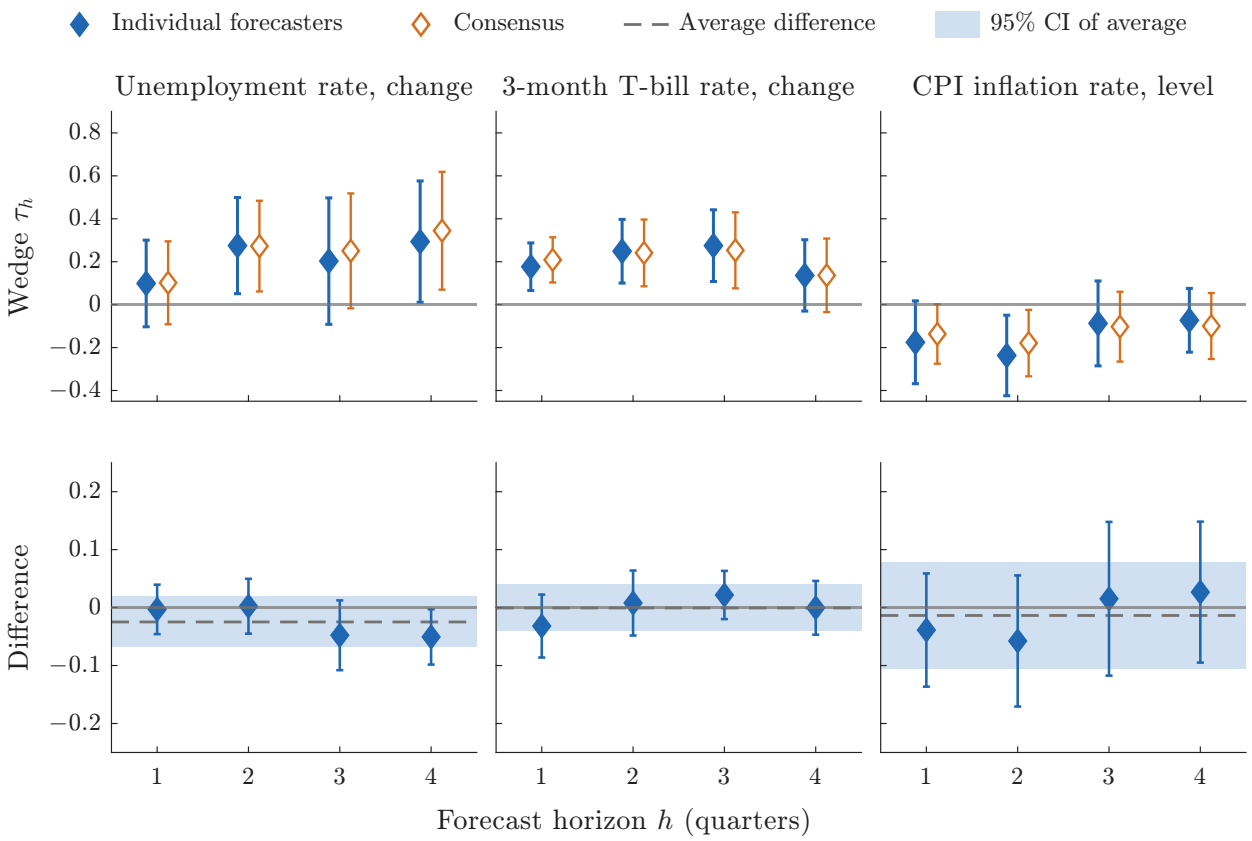


Figure A2. The propagation wedge in the individual SPF responses

Notes: Series, transformations and sources are those of Figure 1, applied to each individual SPF response over forecast origins 1984Q1–2018Q4; in change form the nowcast of the change is the forecaster's own current-quarter value less her own previous-quarter value. Each column is one headline series. The top row plots the propagation wedge τ_h estimated on the individual responses (filled blue diamonds) beside the consensus wedge of Figure 1 (open orange diamonds), each with its 95 percent confidence interval, at horizons $h = 1, \ldots, 4$ quarters. The bottom row plots the individual minus the consensus wedge (filled diamonds) with its 95 percent confidence interval, and its average over the four horizons (dashed line) with the 95 percent band for that average (shaded); its vertical scale is finer than the top row's. The samples are the responses that report a nowcast and all four forecasts: 4,854 from 224 forecasters for the unemployment rate, 4,589 from 219 for the Treasury-bill rate and 4,710 from 225 for inflation. The individual wedge pools the responses with each series centered on the forecaster's own mean; its standard errors are Driscoll–Kraay over forecast origins, with Bartlett weights and four lags, and the difference is estimated jointly with both wedges; Appendix B states the estimator.

end after 2018Q4, through the last origin whose four-quarter outcome has been released, reach the 2020 quarters—as origins when they end in 2020Q1 or later, through their four-quarter outcomes when they end in 2019—and they show why the estimation sample stops before them: the Treasury-bill wedge keeps its sign, but the unemployment wedge turns slightly negative, and the inflation wedge turns positive in the later windows, which contain the 2021–2023 surge. Section 3.5 plots every window and shows that the unemployment reversal follows a reversal of actual propagation, as the mechanism of Section 3 requires.

The window statistics of Figure 7 apply the estimator of Figure 1 window by window, each series centered on its own mean within the window. With $\bar{d}$, $\bar{p}$ and $\bar{q}$ the window means of the products d_t, $p_{t,h}$ and $q_{t,h}$ of Appendix B1, the last two averaged over the four horizons, $\overline{\tau} = (\bar{p} - \bar{q})/\bar{d}$ and $\bar{a} = \bar{p}/\bar{d}$, so the implied weight is $(\bar{p} - \bar{q})/\bar{p}$ in change form and $(\bar{q} - \bar{p})/(\bar{d} - \bar{p})$ in level form; its standard error is Newey–West with four lags, by the delta method, as for $\overline{\tau}$.

The window ratio is biased toward zero in change form: under a constant weight equal to the 1984–2018 estimate, in the data-generating process described next, the median window ratio is 0.41, 0.70 and 0.19 against true weights of 0.52, 0.77 and 0.17, so the text compares window values with window values only; the bands are delta-method, and in change form they are about 40 percent narrower than the sampling spread of the ratio in that process.

That process is the *constant-weight null*: the calibrated rational consensus of Section 2.4 with projection bias at the 1984–2018 weight added (0.52, 0.77 and 0.17), so that $\tau_h = \lambda g_h$ exactly, where g_h is a_h in change form and $-(1 - a_h)$ in level form. In it, 2,000 samples of 140 origins give, across their 81 windows, a correlation between the implied weight and the window standard deviation with median $+0.45$, $+0.29$ and -0.21 and 5th–95th percentiles $[-0.34, 0.85]$, $[-0.38, 0.82]$ and $[-0.86, 0.72]$, against observed correlations of -0.94, $+0.49$ and -0.28; the slope of the wedge on the factor times the standardized volatility, with the factor itself as the other regressor and no intercept, has median $+0.05$, $+0.02$ and -0.01 and percentiles $[-0.11, 0.25]$, $[-0.10, 0.14]$ and $[-0.08, 0.05]$, against ob-

served slopes of -0.14, $+0.03$ and -0.01.

E4. The measured wedge in every cell of the panel

Figure A3 plots the measured average wedge of every cell of the panel of Table 3—each series in each form it supports—with its 95 percent confidence interval. Section 3.4 states the joint tests these cells deliver, and Figure 6 sets the same cells against the wedge one anchor weight implies.

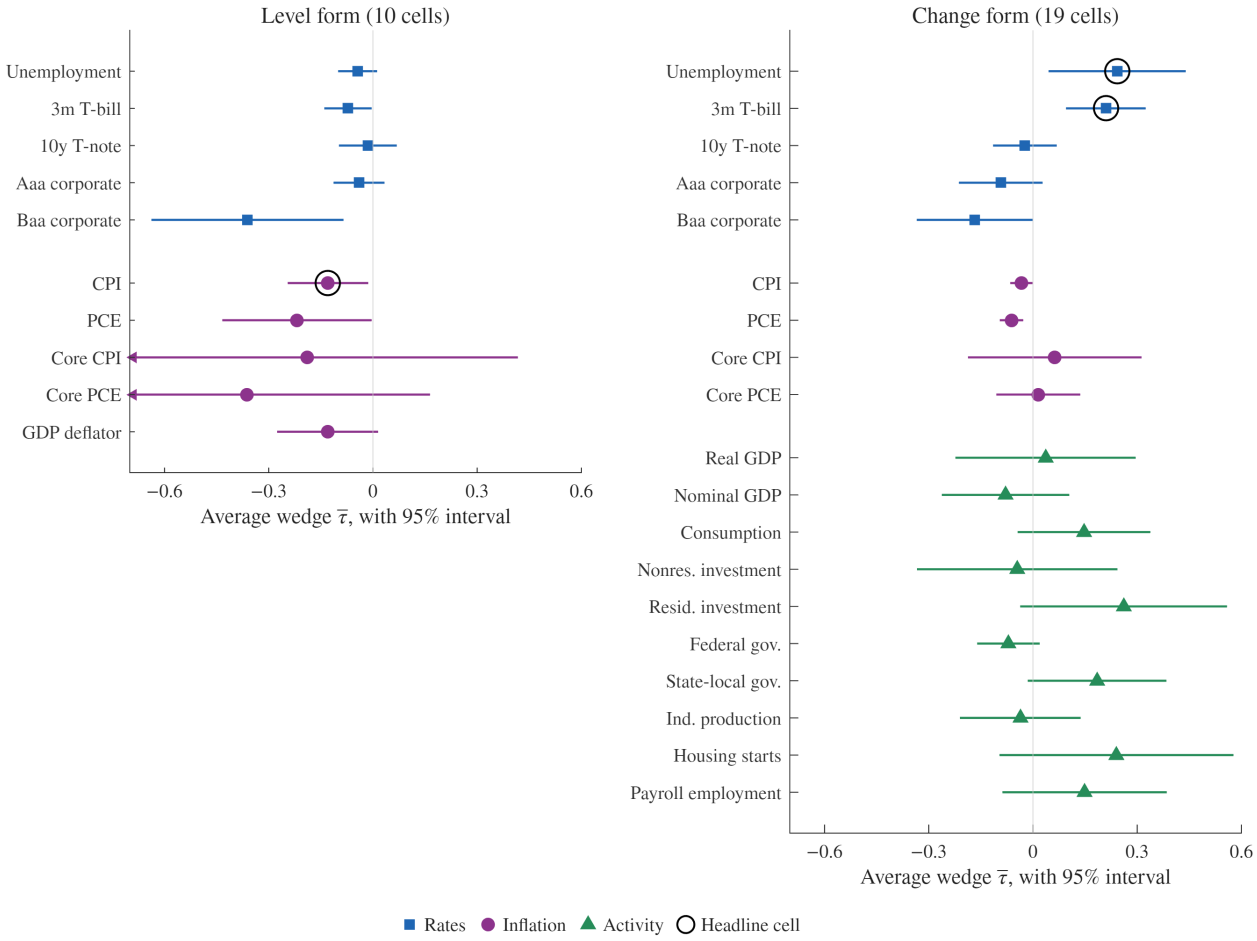


Figure A3. The measured propagation wedge in every cell of the SPF panel

Notes: Series, transformations and samples are those of Table 3, and each wedge is estimated as in Figure 1. Each row is the average wedge $\bar{\tau}$ over $h = 1, \ldots, 4$ in one cell—one series in one form—with its 95 percent confidence interval, and ringed markers are the three headline cells. Rows run by domain and then in the order of Table 3, not by the size of the estimate; the two panels share one row pitch, so the first nine rows, the series that support both forms, line up across them. A whisker that runs off the axis ends in an arrowhead; only core CPI and core PCE in level form, whose standard errors are 0.31 and 0.27, do so. By a moving-block bootstrap over calendar quarters (8-quarter blocks, 1,999 draws, applied to all cells at once), the joint test that all 29 wedges are zero has $p = 0.01$, the test on the 19 change-form wedges alone $p = 0.004$, and the test on the 10 level-form wedges alone $p = 0.26$. All ten level-form wedges are negative.

Table A5 sets the three rival diagnostics of Section 1.2 beside the wedge in the same cells.

Table A5. The three rival diagnostics beside the wedge in every cell of the panel

	Wedge	Error on outcome	Error on nowcast	Persistence comparison
Level form				
Unemployment rate	−0.044	−0.028	−0.046	−0.025
3-month T-bill	−0.072	−0.065	−0.073	−0.064
10-year T-note	−0.015	+0.001†	−0.017	−0.010
Aaa corporate bond	−0.040	−0.058	−0.068	−0.039
Baa corporate bond	−0.361	−0.305	−0.393	−0.363
CPI	−0.130	−0.059	−0.173	−0.357
PCE	−0.219	−0.079	−0.369	−0.186
Core CPI	−0.189	−0.066	−0.332	−0.542
Core PCE	−0.362	−0.036	−0.373	−0.575
GDP deflator	−0.131	−0.044	−0.236	−0.464
Change form				
Unemployment rate	+0.242	+0.287	+0.370	+0.221
3-month T-bill	+0.210	+0.263	+0.267	+0.260
10-year T-note	−0.023	−0.008	−0.050	−0.010
Aaa corporate bond	−0.093	+0.005†	−0.096	−0.010
Baa corporate bond	−0.168	−0.162	−0.243	−0.187
CPI	−0.033	−0.107	−0.104	−0.132
PCE	−0.062	−0.125	−0.212	−0.136
Core CPI	+0.063	−0.144†	+0.028	−0.169†
Core PCE	+0.015	−0.103†	+0.021	−0.135†
Real GDP	+0.036	+0.103	+0.055	−0.054†
Nominal GDP	−0.079	+0.063†	−0.071	−0.192
Consumption	+0.147	+0.138	+0.193	−0.017†
Nonres. investment	−0.045	+0.119†	+0.085†	−0.091
Resid. investment	+0.261	+0.232	+0.288	+0.023
Federal gov.	−0.071	+0.004†	+0.081†	−0.195
State-local gov.	+0.185	+0.114	+0.406	−0.302†
Ind. production	−0.037	+0.131†	+0.046†	+0.045†
Housing starts	+0.240	+0.085	+0.032	−0.209†
Payroll employment	+0.149	+0.272	+0.293	+0.191

Notes: Series, transformations and samples are those of Table 3, and each row is one cell of the panel. The columns report the average propagation wedge $\bar{\tau}$ of Figure A3 and the three rival diagnostics of equation (12) on the same cell, each averaged over $h = 1, \ldots, 4$ and signed as actual minus perceived, so that a positive value reads as under-propagation: error on outcome (the left column of M_h), error on nowcast (the right column) and the persistence comparison (the main diagonal). A dagger marks a rival whose sign differs from the wedge's in that cell. The counts below use 95 percent intervals from the moving-block bootstrap of Figure A3 (8-quarter blocks, 1,999 draws), applied to all cells and all four statistics at once. Across the 29 cells the three rivals agree in sign with the wedge in 21, 26 and 22, and their intervals exclude zero in 10, 10 and 15, against the wedge's 5; all three headline cells agree with every rival.

E5. The level–change test on every series that supports both forms

Table A6 extends the level–change test of Section 3.2 from the three headline series to every series that supports both forms: one weight per series, fitted on its headline form, predicts the average wedge of its other form with no free parameter.

Table A6. The level–change test on every series that supports both forms

	Headline		Counterpart	Average wedge $\overline{\tau}$		p-value	
Series	form	λ [95 percent]	form	measured	implied	$\overline{\tau}=0$	$\overline{\tau}=$ implied
Rates							
Unemployment	change	0.52 [0.28, 0.74]	level	−0.044	−0.050	0.13	0.79
3-month T-bill	change	0.77 [0.49, 0.94]	level	−0.072	−0.083	0.05	0.54
10-year T-note	change	n.i.	level	−0.015	(−)	0.61	—
Aaa corporate bond	change	n.i.	level	−0.040	(−)	0.23	—
Baa corporate bond	change	1.09 [0.21, 2.02]	level	−0.361	−0.522	0.004	0.20
Inflation							
CPI	level	0.17 [−0.02, 0.32]	change	−0.033	−0.015	0.02	0.11
PCE	level	0.20 [−0.13, 0.30]	change	−0.062	−0.022	0.006	0.06
Core CPI	level	n.i.	change	+0.063	(+)	0.52	—
Core PCE	level	n.i.	change	+0.015	(+)	0.81	—

Notes: Series and samples are those of Table 3, and bold marks the headline series of Figure 4. Each series' anchor weight λ is fitted on its headline cell alone—the change form for the rates, the level form for inflation—as in Figure 4: $\lambda = \overline{\tau}/\bar{a}$ in change form, $\lambda = -\overline{\tau}/(1-\bar{a})$ in level form. The implied average wedge of the counterpart cell, $\lambda\bar{a}$ in change form and $-\lambda(1-\bar{a})$ in level form at that cell's own $\bar{a}$, carries no free parameter. Its sign needs only $\lambda > 0$: negative in level form, the sign of the cell's own $\bar{a}$ in change form. The sign alone is shown, in parentheses, where λ is not identified (n.i.): the headline change cells of the 10-year note and Aaa have actual propagation near zero ($\bar{a} = -0.06$ and -0.07), so the ratio divides by nearly nothing, and the core price series start in 2007 ($N = 48$). Baa's weight exceeds one, the bound a projection weight should respect; its interval contains one. Intervals in brackets and p-values are from a moving-block bootstrap over calendar quarters (8-quarter blocks, 4,999 draws) applied to both cells of a series at once; p-values are two-sided, from recentred draws. The measured sign matches the implied one in all nine counterpart cells.

E6. Is the weight one number?

Table A6 fits one weight per series; Figure 6 fits one across the panel. Whether the data distinguish the two is tested here on the moving-block resamples of Figure 6—blocks of eight consecutive quarters applied to all 29 cells at once, 1,999 draws, the same draws behind every bootstrap number in Section 3.4—so that the per-series weights, their differences and the held-out fits are all drawn together. Each per-series weight is the ratio of averages on its own headline cell, the estimator of Figure 4.

The three headline weights differ. Unemployment is 0.52 [0.27, 0.75], the Treasury bill

0.77 [0.49, 0.94] and CPI inflation 0.17 [−0.02, 0.32]; the pairwise differences are −0.25 [−0.52, 0.07] ($p = 0.085$), 0.35 [0.05, 0.64] ($p = 0.031$) and 0.59 [0.30, 0.81] ($p = 0.006$), two-sided from recentred draws, and the joint Wald statistic on the bootstrap covariance is 13.1 on two degrees of freedom, $p = 0.010$ against its own recentred law (0.027 and 0.007 with blocks of four and twelve quarters). Adding the Baa yield and the PCE deflator, whose own weights Table A6 reports, gives $p = 0.024$ on four degrees of freedom. That rejection is not robust to the length of the blocks: with blocks of four, eight and twelve quarters it is $p = 0.076$, 0.024 and 0.014, where the test on the three headline weights gives 0.027, 0.010 and 0.007; the Baa weight is fitted on thirty-six origins from 2010Q1, and with four-quarter blocks its interval, [−0.14, 2.58], contains zero. The rejection is a contrast between groups: the ratio of averages on the two rate cells is 0.61 [0.42, 0.77], on the inflation cell 0.17 [−0.02, 0.32], and their difference 0.44 [0.20, 0.68] ($p = 0.006$).

The held-out cells lean the same way, less sharply. On their own, the eight rate cells outside the fit want 0.71 [0.49, 0.85], the eight price cells 0.38 [−0.05, 0.63] and the ten activity cells 0.19 [−0.13, 0.49]; the held-out rate-minus-price contrast is 0.33 [−0.05, 0.83], $p = 0.10$. Predicting the held-out cells with two weights instead of one—the rate weight for the rate cells, the inflation weight for the price cells—leaves the uncentred R^2 on those sixteen cells at 0.76 against 0.74, a difference whose bootstrap law straddles zero, and matches the same sixteen signs; within the rate cells alone the two-weight fit is 0.92 against 0.73, a gain of 0.18 [0.04, 0.28]. Activity does not belong with the rates: assigning the rate weight to the ten activity cells lowers their fit from 0.24 to −0.83, since they want a weight near the inflation one. Source: `empirics/lambda_equality_bootstrap.py`.

By least squares through the origin, which weights each cell by the square of its factor, the weight of the three headline cells and the held-out weight of Figure 6 are 0.32 and 0.30, and the second has probability 0.08 under no wedge (0.04 with precision weights). With 4- and 12-quarter blocks the held-out weight has probability 0.05 and 0.03 under no wedge, and the joint test of Section 3.4 gives $p = 0.007$ and $p = 0.008$. On the 15 held-out cells with the full 140-quarter window, the weight is 0.24 ($p = 0.06$). Two cells have bootstrap standard errors more than twice their Newey–West counterparts, core PCE in level form and payroll employment in change form, both on samples that start after 2003.

Source: `empirics/wedge_bootstrap.py`.

E7. The anchor and mean squared forecast error

The exercise keeps each series' weight fixed at the value fitted on its headline wedge and asks whether removing the resulting anchor lowers the squared forecast error. It makes that comparison at each of the four horizons and after pooling them, both in sample and in real time; in the real-time exercise, each forecast uses a weight fitted only on outcomes already released by its forecast origin.

The correction produces its only economically material improvement for the unemployment rate in level form: it lowers the pooled mean squared error by 16 percent in sample and 19 percent in real time, although neither gain is statistically significant. For the Treasury bill, removing the wedge-implied anchor raises the pooled mean squared error by more than 200 percent in sample and more than 1,100 percent in real time in both forms. The other cells show no material or statistically significant pooled improvement.

E8. The error–revision check: robustness and the benchmark's sample

Table A8 reports the check of Section 3.3 on the two halves of the sample, with the weight re-fitted on each half, with block-bootstrap intervals on the full sample, and in the four-quarter-average specification of Coibion and Gorodnichenko (2015). Table A9 replicates the benchmark's own regression on the GDP deflator and walks it, one change at a time, to the paper's CPI cell; the split of the sample at 1984 is where its significance goes.

E9. Why the anchor weight cannot be estimated by regressing forecast errors on forecasts

The natural regression for the weight is the forecast error on the forecast's own predicted change,

$$\Delta z_{t+h} - \mathbb{F}_t \Delta z_{t+h} = \alpha + \theta\, \mathbb{F}_t \Delta z_{t+h} + u, \qquad \theta = \frac{\lambda}{1-\lambda}, \tag{E1}$$

exact under the model, with every object observed. It fails on this data for two reasons, neither of which touches the wedge.

First, the equation has two quantities, not three. The dependent variable is the regressor subtracted from the realized change, so the slope is the slope of the realization on the

Table A7. The share of the mean squared forecast error the anchor accounts for

	$h = 1$	$h = 2$	$h = 3$	$h = 4$	Pooled
Unemployment rate, level form ($\lambda = 0.52$)					
In sample, anchor removed	14.1 (14.1)	20.3 (13.8)	17.6 (11.2)	13.2 (9.9)	15.6 (11.0)
Real time, anchor removed	11.4 [0.66]	25.3 [0.40]	23.2 [0.38]	15.0 [0.51]	18.7 [0.45]
Real time, rigidity corrected	6.0 [0.50]	12.9 [0.36]	14.4 [0.27]	—	—
Unemployment rate, change form (headline cell) ($\lambda = 0.52$)					
In sample, anchor removed	0.7 (24.7)	5.4 (12.7)	−5.5 (10.3)	−12.3 (10.1)	−3.8 (12.8)
Real time, anchor removed	−36.4 [0.39]	6.0 [0.78]	−15.2 [0.30]	−33.8 [0.06]	−19.8 [0.29]
Real time, rigidity corrected	2.5 [0.74]	12.4 [0.31]	−3.4 [0.22]	—	—
3-month Treasury-bill rate, level form ($\lambda = 0.77$)					
In sample, anchor removed	−112 (114)	−163 (151)	−203 (184)	−251 (225)	−215 (194)
Real time, anchor removed	−1153 [0.02]	−1256 [0.02]	−1123 [0.01]	−1112 [0.01]	−1140 [0.01]
Real time, rigidity corrected	16.9 [0.09]	16.7 [0.05]	19.3 [0.04]	—	—
3-month Treasury-bill rate, change form (headline cell) ($\lambda = 0.77$)					
In sample, anchor removed	−188 (181)	−178 (164)	−224 (209)	−294 (277)	−222 (205)
Real time, anchor removed	−2023 [0.02]	−950 [0.00]	−775 [0.00]	−956 [0.00]	−1129 [0.01]
Real time, rigidity corrected	5.3 [0.19]	−2.0 [0.37]	−2.2 [0.21]	—	—
CPI inflation rate, level form (headline cell) ($\lambda = 0.17$)					
In sample, anchor removed	0.5 (1.6)	1.0 (1.9)	−1.0 (2.1)	−2.0 (2.6)	−0.4 (1.5)
Real time, anchor removed	−0.0 [0.99]	3.0 [0.43]	−0.3 [0.92]	−2.0 [0.53]	0.2 [0.95]
Real time, rigidity corrected	0.5 [0.62]	−0.4 [0.46]	−0.8 [0.20]	—	—
CPI inflation rate, change form ($\lambda = 0.17$)					
In sample, anchor removed	5.9 (2.7)	0.6 (0.5)	0.1 (0.2)	−0.1 (0.2)	1.3 (0.7)
Real time, anchor removed	8.8 [0.13]	1.6 [0.20]	0.4 [0.34]	−0.2 [0.59]	2.1 [0.14]
Real time, rigidity corrected	9.9 [0.40]	−0.3 [0.68]	−0.3 [0.51]	—	—

Notes: Series, transformations, sample and sources are those of Figure 1, for the six cells of Figure 4. Each entry is the share of the mean squared error of the consensus forecast removed by netting out the anchor, in percent, by horizon and pooled over $h = 1, \dots, 4$; a negative share means the correction raises the error. In the in-sample rows the forecast net of the anchor divides every forecast of a future change by $1 - \lambda$ (in level form, $(\mathbb{F}_t z_{t+h} - \lambda\, \mathbb{F}_t z_t)/(1 - \lambda)$), with the series' weight of Figure 4, fitted on covariances with the present and not on the errors; standard errors, in parentheses in the in-sample rows, are by the delta method, Newey–West with four lags, charging the weight's sampling error. In the real-time rows the weight is re-fitted at each forecast origin from 1994Q1 to 2018Q4 on the origins whose four-quarter outcomes have been released ($s \le t - 5$), the forecast for t is corrected with it, and the share sums squared errors over the 100 out-of-sample origins; the two-sided p-value of the Diebold–Mariano test with the Harvey–Leybourne–Newbold correction and Bartlett lag h (four when pooled) is in brackets in the real-time rows. The rigidity-corrected rows do the same in real time with the forecast moved by the error-on-revision slope of Coibion and Gorodnichenko (2015), re-fitted at each origin on the errors already released, slope only; the same-target revision exists at $h \le 3$.

Table A8. The error–revision check on two halves, with bootstrap intervals, and in the four-quarter-average specification

	$h = 1$	$h = 2$	$h = 3$	Four-quarter average
Unemployment rate, change form				
1984Q1–2001Q2, weight 0.68 (0.19), standard errors				
Observed	+0.42 (0.55)	+1.18 (0.61)	−0.46 (0.63)	—
Net of the anchor	−0.64 (0.24)	−0.19 (0.44)	−0.53 (0.29)	—
2001Q3–2018Q4, weight 0.44 (0.13), standard errors				
Observed	+1.14 (0.70)	+2.42 (0.71)	+3.58 (1.46)	—
Net of the anchor	−0.07 (0.46)	+0.75 (0.36)	+1.50 (0.59)	—
1984Q1–2018Q4, weight 0.52, bootstrap 95 percent intervals				
Observed	+0.89 [−0.32, +1.74]	+1.93 [+0.16, +2.85]	+1.69 [−1.07, +3.72]	+1.12 [+0.12, +1.76]
Net of the anchor	−0.28 [−0.75, +0.26]	+0.38 [−0.32, +0.95]	+0.47 [−0.74, +1.41]	+0.29 [−0.19, +0.71]
3-month Treasury-bill rate, change form				
1984Q1–2001Q2, weight 0.81 (0.16), standard errors				
Observed	+0.72 (0.24)	−0.60 (0.77)	−0.52 (0.78)	—
Net of the anchor	−0.60 (0.19)	−0.52 (0.22)	−0.60 (0.21)	—
2001Q3–2018Q4, weight 0.76 (0.11), standard errors				
Observed	+0.58 (0.30)	+0.07 (0.63)	−0.38 (0.58)	—
Net of the anchor	−0.50 (0.16)	−0.07 (0.19)	−0.13 (0.20)	—
1984Q1–2018Q4, weight 0.77, bootstrap 95 percent intervals				
Observed	+0.67 [+0.22, +1.08]	−0.41 [−1.48, +0.52]	−0.49 [−1.51, +0.45]	+0.75 [+0.37, +1.04]
Net of the anchor	−0.53 [−0.71, −0.20]	−0.28 [−0.79, +0.10]	−0.27 [−0.64, +0.18]	+0.07 [−0.12, +0.26]
CPI inflation rate, level form				
1984Q1–2001Q2, weight 0.18 (0.13), standard errors				
Observed	+0.36 (0.31)	+0.02 (0.83)	+0.12 (0.41)	—
Net of the anchor	+0.89 (0.57)	−0.05 (0.68)	−0.16 (0.45)	—
2001Q3–2018Q4, weight 0.11 (0.09), standard errors				
Observed	+0.35 (0.67)	−1.89 (1.46)	−1.03 (1.42)	—
Net of the anchor	+0.18 (0.85)	−0.84 (1.01)	−0.49 (0.60)	—
1984Q1–2018Q4, weight 0.17, bootstrap 95 percent intervals				
Observed	+0.37 [−0.52, +0.90]	−0.41 [−2.09, +1.04]	−0.12 [−1.33, +0.90]	+0.19 [−0.17, +0.56]
Net of the anchor	+0.44 [−0.58, +1.22]	−0.15 [−1.65, +0.73]	−0.22 [−0.80, +0.64]	+0.21 [−0.31, +0.64]

Notes: Series, transformations, sample and sources are those of Figure 1, and the regression and the forecasts net of the anchor are those of Figure 5. The first two blocks of each series split the 140 forecast origins at 2001Q2: in each half the weight is re-fitted on that half's wedge (its value and standard error in the block label) and both slopes are recomputed, with delta-method Newey–West standard errors in parentheses that charge the weight's error. The third block uses the full sample with a moving-block bootstrap, blocks of 8 quarters over the origins and 1,999 draws; on each draw the weight and both slopes are recomputed jointly, and the interval in brackets is the 2.5th to 97.5th percentile. The last column is the four-quarter-average specification: the error of the average forecast over the current quarter and the next three on the revision of that average between consecutive surveys, the aggregation of Coibion and Gorodnichenko (2015); net of the anchor, each forecast one or more quarters ahead is corrected as in Figure 5 and the nowcast is left as it is; the observed slope carries a Newey–West standard error and both slopes the bootstrap interval.

Table A9. The benchmark regression replicated, and walked to the paper's inflation cell

Specification	β (standard error)	t	R^2	N
Deflator, four-quarter average, 1969–2014 (the benchmark)	+1.209 (0.495)	2.44	0.198	175
Deflator, four-quarter average, 1969–1983	+1.310 (0.535)	2.45	0.223	55
Deflator, four-quarter average, 1984–2014	+0.200 (0.243)	0.82	0.008	120
Deflator, four-quarter average, 1984–2018	+0.185 (0.233)	0.79	0.007	132
CPI, four-quarter average, 1981Q3–2014	+0.291 (0.185)	1.57	0.015	133
CPI, four-quarter average, 1984–2018	+0.171 (0.193)	0.89	0.005	140
CPI, four-quarter average, 1984–2018, Table 3 units	+0.191 (0.192)	1.00	0.006	140
Deflator, $h = 1$, 1969–2014	+0.637 (0.455)	1.40	0.050	183
Deflator, $h = 2$, 1969–2014	+0.928 (0.574)	1.61	0.052	183
Deflator, $h = 3$, 1969–2014	+1.089 (0.505)	2.16	0.070	178
Deflator, $h = 1$, 1984–2018	−0.010 (0.213)	−0.05	0.000	138
Deflator, $h = 2$, 1984–2018	−0.088 (0.494)	−0.18	0.000	138
Deflator, $h = 3$, 1984–2018	+0.135 (0.353)	0.38	0.001	138
CPI, $h = 1$, 1984–2018, Table 3 units (paper cell)	+0.374 (0.340)	1.10	0.005	140
CPI, $h = 2$, 1984–2018, Table 3 units (paper cell)	−0.409 (0.820)	−0.50	0.003	140
CPI, $h = 3$, 1984–2018, Table 3 units (paper cell)	−0.123 (0.577)	−0.21	0.000	140

Notes: Each row is the slope of the consensus forecast error on the same-target forecast revision, with an intercept and a Newey–West standard error with four lags in parentheses, on the origins named. The first row is the benchmark of Coibion and Gorodnichenko (2015): the error of the SPF mean forecast of average GDP-deflator inflation over the current quarter and the next three, in annualized percent as reported, on the revision of that average since the previous survey, with real-time deflator inflation as the outcome. The rows below change one thing at a time—the sample, the deflator for the CPI, the four-quarter average for single-quarter targets, the reported units for the units of Table 3—down to the paper's CPI cell, the last three rows.

forecast less one, $\widehat{\theta} = \widehat{\beta} - 1$, identically, and no rearrangement of the equation changes the estimate. The weight is then $1 - 1/\widehat{\beta}$, which is violently nonlinear near one: a slope of 0.95 gives a weight of -0.05, a slope of 0.5 gives -1, and the ordinary imprecision in a forecast, which pulls an ordinary slope toward zero, becomes an absurd weight.

Second, the regressor is a forecast—predetermined but not strictly exogenous, since later forecasts respond to outcomes that today's error is part of—and it is persistent, so least squares is correct only in large samples, the finite-sample error familiar from dynamic panels and from the prediction of returns by a persistent variable. Simulating an economy in which forecasters are fully rational, so that the true weight is zero, and running equation (E1) on it:

Persistence of the state, sample length	estimated θ	implied weight
0.95, 140 quarters (the sample)	0.600	0.302
0.95, 500 quarters	0.136	0.060
0.95, 5,000 quarters	0.025	0.017
0.50, 140 quarters	0.020	-0.049

At the length of the sample and the persistence of a macroeconomic series, a world with no anchor at all returns a weight of about 0.30; the error is the same with one forecaster and with two hundred, and it shrinks only with the length of the time series.

The regression of Bordalo et al. (2020) is safe from both problems. It regresses the error on the revision of the forecast since the previous survey, a separate number, so there is no collapse to two quantities; and a revision is close to news, far less persistent than a predicted change, so the finite-sample error that is decisive here is small there. Differencing across horizons does not help, since the predicted change still moves with the same slow-moving nowcast; only differencing across time, which a revision does, reduces the persistence.

None of this touches the wedge. Every quantity in it pairs a forecast with a realized outcome; the forecast error never appears, and the forecast never appears on both sides of anything.