Overreaction Through Expectation Smoothing*

Constantin Bürgi† Julio L. Ortiz‡

January 20, 2022

Abstract

We document that updates to macroeconomic expectations among professional forecasters are: (i) negatively serially correlated at the individual level, (ii) positively serially correlated at the aggregate level, and (iii) exhibit an offsetting pattern. To explain these facts, we build and estimate a model featuring annual smoothing and a requirement that quarterly predictions be jointly consistent with the annual forecast. Relative to existing theories, our model provides a unified explanation for these facts as well as other forms of over- and underadjustment. Furthermore, our model suggests that annual forecasts exhibit more information rigidity than quarterly forecasts, with a larger role for sticky information relative to noisy information.

JEL: C53, D83, D84, E17, E27, E37, E47

Keywords: Information Friction, Temporal Consistency, SPF, Inattention, Overreaction.

---

*We would like to thank Philippe Andrade and the participants in the UCD seminar for their valuable comments and support.

†Bürgi (corresponding author): *Email: constantin.burgi@ucd.ie, University College Dublin, Belfield, Dublin 4, Ireland*

‡*Email: julio.l.ortiz@frb.gov*, Federal Reserve Board, 20th and Constitution Avenue NW, Washington DC 20551. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

1 Introduction

Professional forecasters commit predictable mistakes. While individual predictions have been shown to exhibit overreactions (Bordalo et al., 2019, 2020; Broer and Kohlhas, 2019; Bürgi, 2016), aggregate forecasts are characterized by inertia (Coibion and Gorodnichenko, 2015; Dovern et al., 2015). Both forms of error predictability are incompatible with full information rational expectations (FIRE), a benchmark assumption made in macroeconomics. Consequently, theories of non-rational expectations as well as models of imperfect information have been devised to explain over- and underadjustments.[1]

In this paper, we document three key facts relating to survey expectations. First, fixed-event revisions are negatively serially correlated among individual forecasters. Second, fixed-event revisions are positively serially correlated at the consensus-level. Third, forecast revisions across horizons exhibit an offsetting pattern. The third fact, which has not been previously documented in the literature, cannot be reconciled with existing theories of expectation formation. We develop a model of long-run smoothing that can reproduce all of these empirical patterns.

Our model is a version of a hybrid sticky-noisy information model as in Andrade and Bihan (2013) with a focus on the interaction between quarterly and annual forecasts. Two key assumptions are responsible for generating overreactions: temporal consistency (i.e. quarterly forecasts aggregate up to annual forecasts) and stronger smoothing of annual expectations relative to quarterly expectations. With these two assumptions, an upward revision in the near-term must be offset by a downward revision later in the year, as in the data.[2] These overrevisions introduce volatility to quarterly updates which generate overreactions.

We begin by providing empirical evidence relating to overreactions, underreactions, and

[1] Bordalo et al. (2020) for instance present a model of diagnostic expectations while Woodford (2001), Sims (2003), and Mankiw and Reis (2002) present theories of imperfect information.

[2] In the event that forecasts are rounded, quarterly updates would need to be sufficiently large to generate offsetting revisions. At the same time, revisions should not be too large that they lead to a full outlook revision (Baker et al., 2020). While these factors may be present in the data, we nonetheless uncover robust evidence of offsetting revisions, leading us to abstract from rounding and state dependent updating in our model.

annual smoothing based on data from the U.S. Survey of Professional Forecasters (SPF). With regard to overadjustments, we document a robust negative autocorrelation of revisions (Nordhaus, 1987). We repeat this exercise at the aggregate level, documenting a positive autocorrelation of consensus revisions.[3] Finally, we find that when a forecaster revises upward today, she simultaneously revises downward further along her forecasted annual path. We interpret this result as evidence of annual smoothing, and note that existing models of expectation formation cannot flexibly account for offsetting revisions across horizons.

Motivated by these facts, we devise a noisy information model with heterogeneous updating rates by frequency. Forecasters issue quarterly and annual forecasts based on private and public signals. Quarterly and annual updating are separate activities governed by distinct Calvo-like probabilities. Furthermore, forecasters are subject to a consistency constraint which requires that a forecaster's sequence of quarterly predictions aggregate up to her annual forecast.[4]

Excess inattention with respect to annual forecasts can reflect deeper real-world features of professional forecasting. For instance, reputational considerations can generate annual smoothing.[5] Alternatively, time and resource constraints associated with widespread model revisions can generate infrequent annual updating. In both of these settings, forecasters could find it optimal to revise high frequency forecasts while keeping their low frequency outlooks stable.

The source of overreactions in our model comes from forecasters introducing past errors into their reported predictions through the annual consistency constraint.[6] Suppose, for in-

[3]We focus on revision auto-correlation in line with Nordhaus (1987) or Baker et al. (2020), as these over- and underadjustments can be determined ex ante. However, all our results also hold for alternative ex post measures like errors on revisions (Nordhaus, 1987; Coibion and Gorodnichenko, 2015; Bordalo et al., 2020) or errors on actual releases (Kohlhas and Walther, 2021) as shown in the main text and appendix.

[4]While the SPF requires forecasters to produce consistent predictions, Bürgi and Ortiz (2021) document that this is also the case for most forecasters when it is not required. Bürgi and Ortiz (2021) also shows that most forecasts almost immediately reflect the latest data releases.

[5]Ways to model such reputational considerations include adjustment costs as in Kucinskas and Peters (2019) or the game theory framework in Ehrbeck and Waldmann (1996).

[6]Similar to the explanation for an apparent forecast bias at the individual level in Bürgi (2017), the overreaction here is consistent with standard forecasting methods.

stance, that a forecaster periodically makes full updates to her GDP forecast and story. In between these full updates, she replaces the quarterly predictions with actual releases, and then offsets the prediction error so as to ensure annual consistency and preserve her accompanying narrative. These offsetting revisions in turn generate a negative autocorrelation of fixed-event updates since forecasters trade off accuracy with consistency.

Annual smoothing is therefore a key ingredient which allows our model to generate individual overadjustments. While traditional models of forecast smoothing (Scotese, 1994; Woodford, 2001; Mankiw and Reis, 2002; Sims, 2003) can only deliver underreactions, our multi-frequency approach allows us to match quarterly overreactions while preserving aggregate underreactions.

We estimate the model using a minimum distance approach. In particular, we estimate the six parameters of our model by targeting eight micro moments in the panel of real GDP forecasts from the SPF. Our estimated model successfully fits both targeted and non-targeted moments in the data. Overall, our estimates imply that sticky long-run expectations can explain meaningful share of observed overadjustments depending on the testable implication considered. The estimated model can also replicate underreactions in consensus forecasts.

In an effort to quantify the importance of our mechanism relative to other theories, we estimate a version of the model with diagnostic expectations Bordalo et al. (2020).[7] When we add diagnostic expectations to our model, we find that our model explains around 40%-80% of overreactions, depending on the specific metric used. This indicates that annual smoothing is an important contributor to overreactions, alongside other forces.

Finally, we use the model to study information rigidities in survey expectations. Our estimates reveal that information frictions differ across frequencies and are more pervasive at the annual level. This might help explain why Andrade and Bihan (2013) find different levels of inattention for annual predictions depending on the specific method used. When averaging across the two frequencies, we recover implied information frictions similar to pre-

[7]Several alternative theories of non-rational expectations can explain overreactive behavior (Daniel et al., 1998; Broer and Kohlhas, 2019). At the same time, overreactions can arise through optimizing behavior subject to attention or memory constraints (Kohlhas and Walther, 2021; Azeredo da Silvera et al., 2020).

vious estimates documented in the literature (Coibion and Gorodnichenko, 2015; Ryngaert, 2017). In addition, the interplay between the frequencies in our model allows us to decompose the friction into noisy and sticky information. We find that noisy information is the dominant source of information frictions at the quarterly frequency while sticky information is the driver of information frictions at the annual frequency.

The rest of the paper is organized as follows. Section 2 documents motivating empirical evidence relating to overadjustments, underadjustments, and annual smoothing. Section 3 presents the hybrid sticky-noisy information model with differential rates of updating. Section 4 discusses the estimation approach and the estimated parameters. Section 5 quantifies the extent to which long-run rigidity can explain short-run overadjustments. Section 6 discusses the implications for information frictions. Section 7 concludes.

2 Facts About Over- and Underadjustments

We first document some empirical facts about professional forecasts. The patterns that we highlight in the data serve as motivating evidence for the model introduced in the subsequent section. Furthermore, we revisit some of these moments when assessing the estimated model's ability to explain observed overadjustments.

The data that we use come from the SPF, a quarterly survey managed by the Federal Reserve Bank of Philadelphia. The survey began in the fourth quarter of 1968, and provides forecasts from several forecasters across a range of macroeconomic variables over many horizons, h. The SPF reports current-year annual predictions which the survey requires to be consistent with the averages of the quarterly forecasts. In this sense, the consistency constraint that we impose in our model is directly motivated by the data.

2.1 Individual Overreactions

First, professional forecasters exhibit overreactive behavior. To show this, we run two sets of panel regressions: revisions on past revisions and errors on revisions.[8] Both regressions were first introduced as a test of weak efficiency in Nordhaus (1987). Let x_{t+h} denote real GDP growth at time $t+h$. Furthermore, let $F_{it}(x_{t+h})$ denote forecaster i's prediction for x_{t+h} devised at time t. With this notation defined, the revision autocorrelation regression is:

$$F_{it}(x_{t+h}) - F_{it-1}(x_{t+h}) = \gamma_i + \gamma_{1,h}\left[F_{it-1}(x_{t+h}) - F_{it-2}(x_{t+h})\right] + \varepsilon_{it+h}. \tag{1}$$

In words, we focus on a fixed-event and project the current forecast revision on its previous value. We are interested in the coefficient in front of the lagged revision, $\gamma_{1,h}$. A negative value of $\gamma_{1,h}$ indicates that an upward forecast revision yesterday predicts a downward forecast revision today.

Table 1 reports the results across three horizons which imply that forecasters overrevise their predictions. For current quarter forecasts, a one percentage point upward revision today predicts a 0.17 percentage point downward revision tomorrow. Forecasters tend to overrevise more strongly at the one- and two-step ahead horizons, with point estimates hovering at around -0.30.

The second set of results, also reported in Table 1, relate to errors-on-revisions. We run the following regression:

$$x_{t+h} - F_{it}(x_{t+h}) = \beta_i + \beta_{1,h}\left[F_{it}(x_{t+h}) - F_{it-1}(x_{t+h})\right] + \epsilon_{it+h}. \tag{2}$$

When $\beta_{1,h} < 0$, an upward revision predicts a more negative subsequent forecast error, implying that forecasters overreact to new information when updating their predictions. Table 1 reproduces these estimates in our sample. Across horizons, we find that a one percentage point upward forecast revision predicts a roughly -0.23 to -0.30 percentage point more neg-

[8] We provide an additional estimate of overreactions based on Kohlhas and Walther (2021) in Appendix A.

Table 1: Overreaction Among Individual Forecasters

	$h = 0$		$h = 1$		$h = 2$	
	Revision	Error	Revision	Error	Revision	Error
Previous revision	-0.174***		-0.264***		-0.359***	
	(0.036)		(0.036)		(0.054)	
Revision		-0.249***		-0.228***		-0.302***
		(0.071)		(0.066)		(0.051)
Forecasters	254	283	254	254	253	250
Observations	4,752	5,859	4,736	4,811	4,494	4,690

Note: The table reports panel regression results from SPF forecasts of real GDP growth based on regressions (1) and (2). Each set of columns refers to a different horizon, from the current quarter to two quarters ahead. Forecaster fixed effects are specified in each regression, and Driscoll and Kraay (1998) standard errors are reported in parentheses. *** denotes 1% significance, ** denotes 5% significance, and * denotes 10% significance.

ative subsequent forecast error. These estimates are broadly in line with the estimates in Bordalo et al. (2020) and Bürgi (2016).

2.2 Aggregate Underreactions

Whereas individual forecasters appear to overreact, consensus predictions exhibit underadjustments. Such aggregate inertia has been an empirical moment of interest to the literature studying information rigidities. Table 2 reports the results based on the time series analogs of (1) and (2), where, instead of specifying individual forecasts, we focus on consensus forecasts.

The estimates in Table 2 provide some evidence of underadjustments at the aggregate-level. For instance based on a simple model of noisy information, the estimated errors-on-revisions coefficient at the one-quarter ahead horizon implies that forecasters place a weight of $1-\frac{1}{1+0.685} \approx 0.59$ on their prior when updating their prediction.[9] These estimated underreactions are consistent with Coibion and Gorodnichenko (2012), Coibion and Gorodnichenko (2015), Bürgi (2016), Baker et al. (2020), and Bordalo et al. (2020).

[9] As stated in Coibion and Gorodnichenko (2015), the Kalman gain arising from such a model is $\kappa = \frac{1}{1+\beta_{1,h}}$.

Table 2: Underreaction in Consensus Forecasts

	$h=0$		$h=1$		$h=2$	
	Revision	Error	Revision	Error	Revision	Error
Previous revision	0.360***		0.380***		-0.011	
	(0.106)		(0.095)		(0.074)	
Revision		0.354**		0.685**		0.667*
		(0.179)		(0.313)		(0.385)
Observations	200	201	200	200	195	200

Note: The table reports panel regression results from SPF forecasts of real GDP growth based on regressions (1) and (2). Each set of columns refers to a different horizon, from zero steps ahead (current quarter) to two steps ahead. Newey-West standard errors are reported in parentheses. *** denotes 1% significance, ** denotes 5% significance, and * denotes 10% significance.

2.3 Offsetting Revisions

Having documented simultaneous over- and underreactions in the survey data, we turn to providing motivating evidence for our mechanism. If forecasters have a tendency to smooth their annual predictions, then multi-horizon revisions should exhibit an offsetting pattern. For instance, if a forecaster receives positive news today, then she will wish to revise her forecast upward. However, if she is inattentive to her annual forecast (or otherwise wishes to smooth it), then she will have to revise upward subject to the adding up constraint. In order for her newly-issued quarterly predictions to reflect her unchanged annual outlook, the upward revision today must be offset by a downward revision elsewhere along her predicted path.

Directly testing for annual smoothing in the data by running Coibion and Gorodnichenko (2015) regressions with annual forecasts and comparing them to the quarterly estimates is not readily feasible for a number of reasons. First, the annual is constructed using the average of the quarterly predictions (in levels) that result in the annual potentially appearing smoother than without this additional channel. Rounding can exacerbate this effect. Second, the horizon of the annual prediction changes with every survey and reduces the sample substantially. For this reason, we devise a model in which annual smoothing is the mechanism through which offsetting revisions arise, and demonstrate that such a model can better fit

Table 3: Offsetting Revisions Among Individual Forecasters

	Three-quarter ahead revision
Two-quarter ahead revision	0.332***
	(0.025)
One-quarter ahead revision	0.083***
	(0.011)
Current-quarter revision	-0.019**
	(0.008)
Fixed effects	Forecaster, Variable, Time
Observations	65,913

Note: The table reports panel regression results from SPF forecasts based on regression (3). Forecaster, macro variable, and time fixed effects are specified, and Driscoll and Kraay (1998) standard errors are reported in parentheses. *** denotes 1% significance, ** denotes 5% significance, and * denotes 10% significance.

the data than a model without annual smoothing.[10]

We provide two sets of results lending support to the notion that forecasters offset their revisions. First, pooling across a range of macroeconomic variables, we show that an upward revision today predicts a contemporaneous downward revision to the three-quarter ahead forecast. Second, we focus on real GDP forecasts and demonstrate that current-quarter revisions move in the same direction as a surprise (proxied by data revisions to GDP) while three-quarter ahead revisions move in the opposite direction.

We begin by regressing the three-quarter ahead revision devised at time t on the current-quarter revision, controlling for the one- and two-quarter ahead revisions as well as forecaster (i), macroeconomic variable (j), and time (t) fixed effects:

$$F_{ijt}(x_{jt+h+3})-F_{ijt-1}(x_{jt+h+3}) = \alpha_i+\alpha_j+\alpha_t+\sum_{k=0}^{2}\alpha_k\left[F_{ijt}(x_{jt+h+k})-F_{ijt-1}(x_{jt+h+k})\right]+\nu_{ijt} \quad (3)$$

The results are reported in Table 3.

Across a range of variables, and controlling for unobserved heterogeneity at the forecaster, variable, and time levels, we find that the current-quarter revision covaries negatively with

[10]In Appendix D, we also consider an alternative driver of offsetting revisions such as a richer driving process, with little qualitative effect on our results.

the three-quarter ahead revision. In traditional rational expectations models, forecasters optimally update their predicted trajectories. Amid positive news, forecasters would revise their trajectories upward, with the magnitude of the revision over longer horizons regulated by the persistence of the driving process. Importantly, such models do not accommodate for the offsetting pattern suggested by the results in Table 3.

Next, we narrow our focus to real GDP forecasts and dig deeper by examining exogenous surprises. In particular, we analyze the response of forecast revisions to a surprise in the target variable, proxied by statistical data revisions. Macroeconomic variables are subject to frequent data revisions that are made by statistical agencies. We construct a series of real GDP data revisions by computing the difference across vintages: $d_t = x_t^{\text{new}} - x_t^{\text{old}}$. For each horizon, we regress forecast revisions devised at time t on realized data revisions observed at time t, controlling for forecaster fixed effects:

$$F_{it}(x_{t+h}) - F_{it-1}(x_{t+h}) = \alpha_i + \alpha_1 d_t + \varepsilon_{it}. \tag{4}$$

Figure 1 plots the point estimates across horizons, with 95% confidence intervals. The estimates indicate that an upward revision to real GDP induces forecasters to revise their current-quarter predictions upward and concurrently revise their three-quarter ahead predictions downward. This figure accords with the estimates reported in Table 3, and indicate that forecast revisions exhibit an offsetting behavior indicative of long-horizon smoothing.

Taken together, professional forecasts exhibit overreactions and inertia. Furthermore, forecasters appear to offset near-term revisions over their longer-term trajectories. We argue that this latter finding can explain some of the observed overrevisions in the data, and explicitly model offsetting revisions in next section.

3 A Model of Offsetting Revisions

We begin by detailing our hybrid sticky-noisy information model. Our model is inspired by Andrade and Bihan (2013) and features annual and quarterly forecasts, each subject to

Figure 1: Effect of Data Revisions on Forecast Revisions

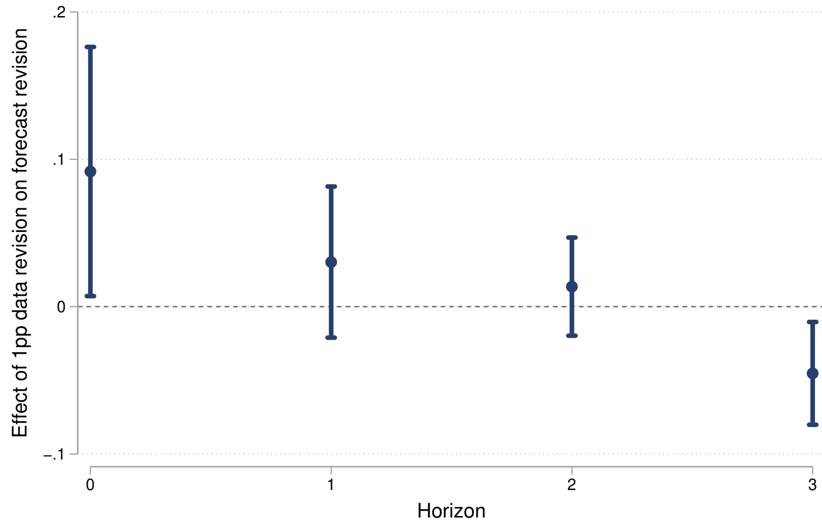

Note: The figure reports 95% confidence estimates of the α_1 coefficient in regression (4) across four horizons. Driscoll and Kraay (1998) standard errors are specified in the regressions.

a distinct updating probability. Derivations of our results can be found in Appendix B. After outlining the model, we discuss how overreactions arise through annual smoothing and temporal consistency. Finally, we analyze a series of comparative statics in order to examine the ways in which the regression coefficients estimated in the previous section depend on the model parameters.

3.1 Model Setup

The model is populated by professional forecasters. Forecasters issue predictions about a macroeconomic variable, which in part reflects the latent state of the economy, subject to the realization of noisy signals. Forecasters issue both quarterly and annual forecasts which they differentially update subject to an adding up constraint which requires that quarterly forecasts jointly respect the annual forecast in every period.

More formally, forecasters aim to predict a macroeconomic variable x_t, which is defined as a function of two components:

$$x_t = s_t + e_t, \quad e_t \sim N(0, \sigma_e^2).$$

The underlying state of the economy, s_t, follows an AR(1) process:

$$s_t = (1-\rho)\mu + \rho s_{t-1} + w_t, \quad w_t \sim N(0, \sigma_w^2)$$

with unconditional mean μ, persistence ρ, and variance $\frac{\sigma_w^2}{1-\rho^2}$. The state of the economy is unobserved to forecasters and to the econometrician. The transitory component, e_t, is normally distributed noise with variance σ_e^2.

Forecasters are interested in forecasting the quarterly and annual realizations of the macroeconomic variable, x_t. Forecaster i's quarterly k-step ahead forecast devised at time t is $\widehat{x}^i_{t+k|t}$. Her annual forecast devised at time t is $\frac{1}{4}\sum_{h=0}^{3} \widehat{x}^i_{t+h|t}$.[11]

When updating their predictions, forecasters observe the previous realization of the macro variable, x_{t-1}, as well as a contemporaneous private signal:

$$y_t^i = s_t + v_t^i, \quad v_t^i \overset{\text{i.i.d.}}{\sim} N(0, \sigma_v^2)$$

In this linear Gaussian set up, forecasters can employ the Kalman filter to determine the optimal forecast, consistent with the conditional expectation. However, forecasters cannot flexibly update their forecasts every period. Instead, in a given period, a forecaster is only able to revise her quarterly prediction with probability q, and annual outlook with probability p.

Infrequent annual updating ($p < 1$) can be motivated by institutional, reputational, or economic considerations. Anecdotally, forecasting institutions avoid revising their annual figures in each month or quarter, opting instead to implement infrequent large model revisions

[11]In general, a forecaster may be interested in predicting some realization of x over an arbitrary horizon, H. In this case, the forecaster will be interested in forecasting $\frac{1}{H+1}\sum_{h=0}^{H} x_{t+h}$.

a couple of times per year. A probability $p < 1$ can be attributed to time or resource constraints associated with undertaking such large model revisions. Alternatively, more rigid annual updating can reflect the value in sticking to a particular "story" to narrate to clients rather than revising in different directions each period. For our purposes, all of these explanations are embedded in the probability p.

The Calvo-like probabilities, q and p, give rise to four distinct cases:

Case 1: With probability $(1-q)(1-p)$, the forecaster does not update at all.

Case 2: With probability $q(1-p)$, the forecaster updates the quarterly forecast, but not the annual. In this case, she updates the quarterly forecast based on the signals received and subject to an adding up constraint.

Case 3: With probability $(1-q)p$, the forecaster updates her annual forecast, but not the quarterly. We interpret this case as a scenario in which the forecaster simply "brings in" the latest macro release x_{t-1} and updates her annual prediction accordingly. Importantly, the forecaster does not touch the rest of the projected quarterly forecasts.[12]

Case 4: With probability pq, the forecaster can update both types of forecasts. These forecasts are both updated based on the two signals they receive, and constitute an overall revision of the forecast.

3.2 Quarterly Overreactions

From the perspective of the model, quarterly overreactions are due to Case 2 forecasting. As a result, the probability $q(1-p)$ will govern the sign and magnitudes of the coefficients reported in Table 1. For general forms of long-run smoothing, the reported Case 2 prediction is:

$$\widehat{x}^i_{t+k'|t+k} = \mathbb{E}_{it+k}(x_{t+k'}) + \frac{1}{H+1}\sum_{h=0}^{H}\Big[\mathbb{E}_{i,t+k-j}(x_{t+h}) - \mathbb{E}_{it+k}(x_{t+h})\Big], \tag{5}$$

where $\widehat{x}^i_{t+k'|t+k}$ denotes forecaster i's reported forecast in period $t+k$ for some future

[12]This scenario does not play an important role in our findings. The estimated baseline model in the next section implies that Case 3 occurs only 4% of the time. In addition, a version of this model which assumes flexible quarterly updating, $q = 1$, delivers similar conclusions.

period, $t+k'$. The subscript $t+k-j$ refers to period in which the long-run forecast was last updated. Finally, $H+1$ refers to the length of the horizon over which forecasts are smoothed. The reported forecast is the sum of the optimal conditional expectation and a term capturing the gap between the path of the outdated annual forecast and what it should be based on the latest information.

Because our central focus is on quarterly and annual updating, we set the relevant horizon length to be $H=3$. Note, however, that as $H \to \infty$, the second term in (5) vanishes and the reported forecast converges to the conditional expectation. This is intuitive: as the horizon over which a forecaster smooths her forecasts expands, the forecaster has more degrees of freedom along which to adjust the trajectory in order to preserve temporal consistency. As a result, she is more flexibly able to report a prediction that is consistent with the optimal forecast.

We can rearrange (5) in order to more transparently characterize the source of overreactions:

$$\widehat{x}^i_{t+k'|t+k} = \underbrace{\frac{3}{4}\mathbb{E}_{it+k}(x_{t+k'}) + \frac{1}{4}\mathbb{E}_{t+k-j}(x_{t+k'})}_{\text{Traditional smoothing motive}} + \underbrace{\frac{1}{4}\sum_{h\neq k'}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - \mathbb{E}_{t+k}(x_{t+h})\right]}_{\text{Source of overreactions}}.$$

The first two terms on the right-hand side of the above expression reflect averaging between current and past forecasts that arises in traditional forecast smoothing models. The last term is responsible for generating overreactions in our model. This sum reflects the differences in the conditional expectations between $t+k$ and $t+k-j$ for the other quarters that comprise the annual path. As current-year events unfold, this sum incorporates past forecast errors.

To see this, note that (5) can be re-written as:

$$\widehat{x}^i_{t+k'|t+k} = \mathbb{E}_{it+k}(x_{t+k'}) + \frac{1}{4}\sum_{h=0}^{k-1}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - x_{t+h}\right] + \frac{1}{4}\sum_{h=k}^{3}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - \mathbb{E}_{it+k}(x_{t+h})\right],$$

where the second term now reflects relevant past forecast errors.

Overreactions arise because annual inattention and temporal consistency introduce past mistakes into the reported prediction. Suppose, for simplicity, that forecasters last updated their predictions in the previous period so that $j = 1$. Then, the above expression becomes:

$$\widehat{x}^i_{t+k'|t+k} = \mathbb{E}_{it+k}(x_{t+k'}) + \frac{1}{4}\big[\mathbb{E}_{it+k-1}(x_{t+k-1}) - x_{t+k-1}\big] + \frac{1}{4}\sum_{h=k}^{3}\big[\mathbb{E}_{it+k-j}(x_{t+h}) - \mathbb{E}_{it+k}(x_{t+h})\big].$$

Based on the second term, if x_{t+k-1} comes in higher than expected, then forecasters will place downward pressure on their current forecast in order to preserve consistency. As a result, a positive rational expectations error today predicts a positive ex-post forecast error tomorrow. These excessive overrevisions are later corrected as new and relevant information arrives in the next period, generating negatively autocorrelated revisions. The trade-off between accuracy and consistency is therefore responsible for producing overreactions in our model.[13]

Figure 2 highlights the key distinction between a model with annual smoothing and temporal consistency relative to a traditional model of imperfect information. The figure plots regression coefficients estimated off simulated data in the two models based on (3). According to a traditional imperfect information model, the three-quarter ahead forecast revision is positively related to the current-quarter revision. In addition, controlling for the current-quarter revision, one- and two-quarter ahead revisions hold no predictive power over the three-quarter ahead revision. Intuitively, in a model with Bayesian updating, the k-quarter ahead revision is equal to the current-quarter revision, multiplied by the persistence of the driving process, raised to the k-th power.[14]

On the other hand, with annual smoothing, the model is able to generate a negative relation between the three-quarter ahead revision and the current-quarter revisions, and a positive relation between the three-quarter ahead revision and the one- and two-quarter ahead revisions. This offsetting pattern, which our model is able to generate, will be respon-

[13] This mechanism is able to generate initial underreaction and "delayed" overreaction, a feature of the data explored in Angeletos et al. (2020) and Bianchi et al. (2021b). See Appendix B for additional details.

[14] Technically, based on the no annual smoothing model, there is perfect multicollinearity as two of the regressors are a linear combination of the third.

Figure 2: Offsetting with Annual Smoothing

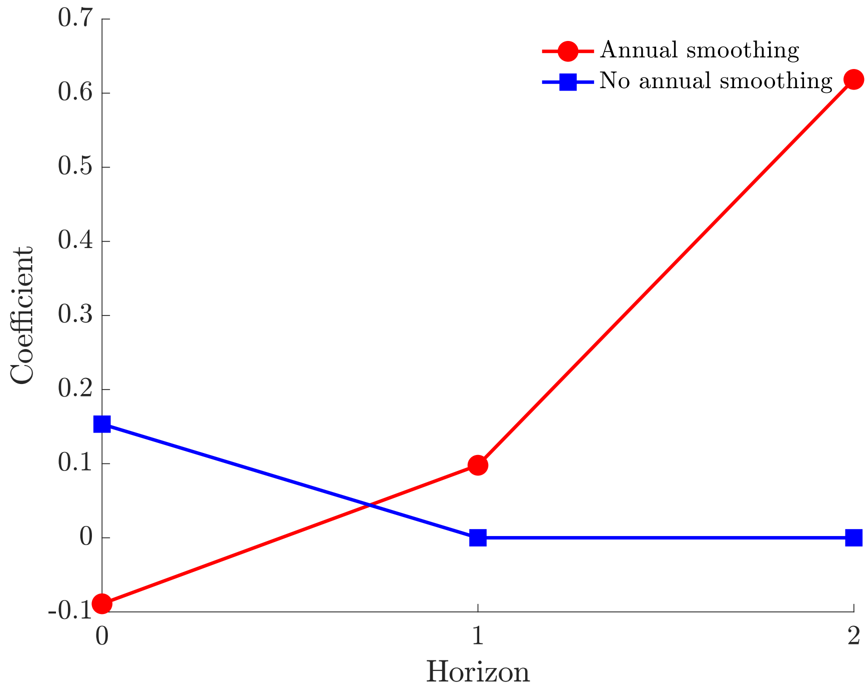

Note: The figure plots the estimated coefficients from simulated regressions based on (3). The red line denotes the model with annual smoothing and the blue line denotes a model without annual smoothing.

sible for producing quarterly overreactions.

3.3 Analyzing the Model

The model features rich dynamics across horizon, frequency, and level of aggregation. As a result, the coefficients studied in Section 2 are complex functions of the underlying model parameters. To extract intuition from the model, we therefore focus on simulated comparative statics.

We focus on the autocorrelation of revisions, though we note that in the same qualitative findings arise when simulating the errors-on-revisions coefficient. Figure 3 plots these coefficients across a range of different parameter values collectively governing the state and signals. For each panel, the left axis plots the individual-level coefficient while the right axis plots the consensus-level coefficient.

Panel 1 displays results for the persistence of the state. As the persistence of the state

Figure 3: Revision Autocorrelation Model Parameters

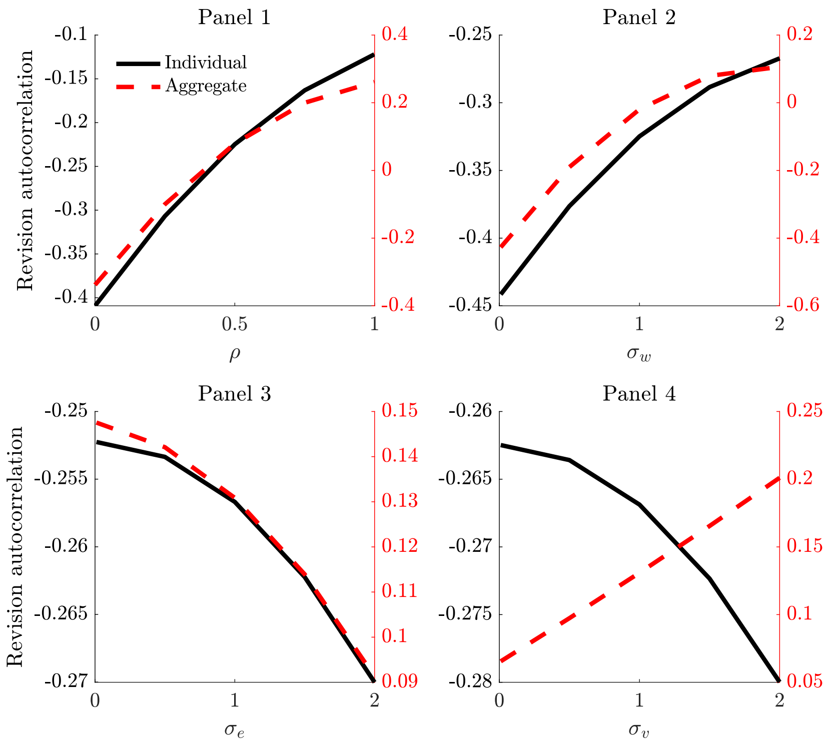

Note: The figure plots comparative statics of one-quarter ahead revision autocorrelation to each model parameter. The black line (left axis) plots the individual pooled autocorrelation of revisions while the dashed red line (right axis) plots the consensus-level autocorrelation coefficient.

approaches one, we find that the scope for overreactions declines. This is consistent with Bordalo et al. (2020) and Afrouzi et al. (2021) who note that overreactions are decreasing in ρ. From the lens of our model, a more persistent target variable will reduce the magnitude of the forecast errors thereby reducing the scope for past forecast errors to influence current predictions through the consistency constraint. Furthermore, the aggregate autocorrelation of revisions rises as the persistence of the state rises. At the consensus-level, the imperfect information environment generates greater persistence in aggregate beliefs with a more persistent driving process.

Panel 2 reports the results for the state volatility, σ_w. Here, we find that the scope for overreactions is decreasing in state innovation volatility. Consistent with the intuition discussed for Panel 1, the rate of learning is increasing in the variance of the latent state. As information becomes more precise, there is less scope of overreactions. We also find that the autocorrelation coefficient increases at the aggregate level as the volatility of the state increases. We uncover similar results when simulating the consensus errors-on-revisions regression coefficient. This finding is at odds with Coibion and Gorodnichenko (2015) which finds that a larger coefficient at the consensus-level is indicative of greater information frictions. Here, the opposite is the case since the Kalman gain is increasing in the volatility of the state.

On the other hand, Panels 3 and 4 show that the forecaster-level coefficients are decreasing in public and private noise. This is because, at the individual level, additional noise raises the variance of the forecast error and reduces the rate of learning. As a result, annual forecast smoothing promotes overadjustments. At the aggregate level, however, the autocorrelation of revisions depends on the type of noise. In particular, the aggregate autocorrelation coefficient falls as common noise becomes more pervasive.[15]. Private noise however, washes out in the cross-section, so the aggregate autocorrelation coefficient actually rises with elevated levels of σ_v. The standard noisy information logic applies here: higher private noise variance reduces the signal-to-noise ratio and the Kalman gains thereby generating inertia in expectation formation.

3.4 Updating Probabilities

An important feature of our model is the differential rates of updating for quarterly and annual forecasts. To assess the role that infrequent annual updating plays in driving observed overreactions, we focus on the frequency of Case 2 forecasting.

Figure 4 illustrates how individual overadjustments depend on the quarterly and annual updating probabilities. On the left axis, the figure plots simulated estimates of the autocorrelation of forecast revisions. The right axis plots the aggregate revision autocorrelation

[15]This is consistent with the discussion in Coibion and Gorodnichenko (2015) on the bias in the OLS errors-on-revisions coefficient under a common noise assumption.

Figure 4: Revision Autocorrelation and Updating Probabilities

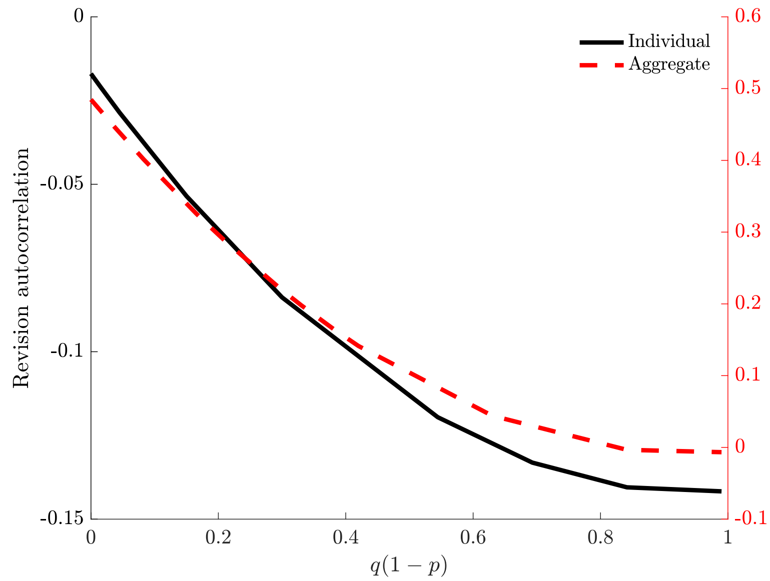

Note: The figure plots the simulated revision autocorrelation coefficients as a function of the probability of Case 2 forecasting. The left axis plots the individual-level coefficient while the right axis plots the aggregate coefficient.

coefficient. Finally, the horizontal axis plots the probability of Case 2 updating.

Focusing first individual autocorrelation coefficient, we see that as the probability of Case 2 updating rises, forecasters' quarterly predictions more intensely overreact. This is because forecasters increasingly find themselves in a scenario in which they wish to update based on news they receive, but cannot adjust the annual outlook. In this case, forecasters respond to news, but offset their sequence of revisions so as to preserve temporal consistency. The excessive revising that occurs along the annual path is responsible for generating overreactions. Turning to the right axis, we note that the simulated aggregate autocorrelation coefficient is also decreasing in the probability of Case 2 updating. As more forecasters engage in Case 2 updating, the role imperfect information in generating inertia diminishes.

4 Model Estimation

Having analyzed the model's ability to reproduce over- and underreactions, we next turn to estimating the model with micro data from the SPF. We fit the model to real GDP growth forecasts between 1968Q4-2019Q4. Of the seven parameters, we first fix the unconditional mean, $\mu = 2.4$, consistent with the sample mean of real-time real GDP growth over this period.

We estimate the remaining six parameters via a minimum distance estimation approach. The parameters to be estimated are $\theta = (\rho \;\; \sigma_w \;\; \sigma_e \;\; \sigma_v \;\; q \;\; p)'$. These parameters are chosen to match eight data moments: the covariance matrix of current-quarter and current-year forecasts, the covariance matrix of current-quarter forecast revisions and last period's forecast error, and the mean squared errors associated with current quarter predictions and current year predictions.

4.1 Identification

As with any other estimation approach, a discussion of identification is imperative. Here, there is a joint mapping between parameters and moments, however, some moments are especially important for identifying certain parameters. Figure C2 illustrates some important comparative statics that lend support to the choice of target moments.

The underlying persistence of the latent state, ρ, is in part identified by the covariance between the current-quarter forecast and the current-year forecast. With a highly persistent data generating process, the covariance between current-quarter and current-year forecasts will be strongly positive. Instead, if the process is i.i.d., then this covariance will be closer to zero. This is seen in Panel 1 of Figure C2. Moreover, Panels 5 and 6 of the figure confirm that the probabilities of updating, q and p, inform the relevant mean squared errors.

The dispersion parameters, σ_w, σ_e, and σ_v require further discussion. Two of these parameters reflect noise variance (σ_e and σ_v) while the other (σ_w) reflects the variance of the underlying state. The distinction between noise and signal is crucial in the identification of these parameters.

First, the variance of the underlying state, σ_w, is identified in part from the variance of the current-year forecast. Recall that the current-year forecast is: $\frac{1}{4}\sum_{h=0}^{3}\widehat{x}^i_{t+h|t}$. As the end of the year approaches, more and more realizations of x_t within the year figure into the optimal current-year projection, replacing the filtered forecasts that are subject to noise. For this reason, an increase in σ_w raises the variance of the current-year forecast.

Moreover, elevated levels of public signal noise, σ_e, contribute to a larger forecast error variance. The link between common noise and the variance of errors is intuitive since the macro variable being predicted is a linear function of the common noise, e_t.

Lastly, private noise variance, σ_v, informs the covariance between revisions and lagged errors. Based on the model, the filtered current-quarter forecast revision is:

$$x^i_{t|t} - x^i_{t|t-1} = \kappa_1(y^i_t - x^i_{t|t-1}) + \kappa_2(x_{t-1} - x^i_{t-1|t-1}).$$

where κ_1 and κ_2 denote the Kalman gains. An increase in σ_v reduces the Kalman gain weight placed on the private signal, κ_1. As σ_v rises, fluctuations in the current-quarter revision are increasingly driven by lagged forecast errors, thereby strengthening the covariance between the revision and the lagged error. In other words, with less informative private signals, forecasters trust y^i_t less and instead base more of their revisions on the news gleaned from yesterday's error.[16]

4.2 Estimation Results

The parameters obtained via the minimum distance estimation approach are precisely estimated and are reported in Panel A of Table 4. The underlying persistence of real-time GDP growth is estimated to be about 0.54. In addition, the dispersion in state innovations is 2.37 while the dispersion in public and private noise are 1.66 and 0.72, respectively. These

[16]Figure C3 in Appendix C helps assess the sources of identification by reporting the sensitivity of each of the six parameters to changes in a given moment, based on Andrews et al. (2017). These figures confirm the intuition laid out above.

estimates imply a signal-to-noise ratio of about $\frac{\sigma_w^2}{\sigma_e^2+\sigma_v^2} \approx 1.72$. Furthermore, the probability of quarterly updating is about 0.93, indicating that forecasters update their quarterly predictions nearly every period. Lastly, the probability of annual updating is estimated to be 0.61, meaning that forecasters update their annual predictions for at most three out of four quarters. Put another way, in any given quarter, roughly 61% of forecasters update their annual forecasts. These estimates are significantly below one, indicating that there is scope for the model to generate both over- and underadjustments. Our estimates imply that annual smoothing is a meaningful friction in the model. In the absence of infrequent annual updating, the root mean squared error for current-quarter predictions would fall by about 10%.

The model is able to successfully replicate the targeted features of the data. Panel B of Table 4 reports the model-implied moments and the empirical moments, scaled to correlations and standard deviations. The fourth column of Panel B reports t-statistics which indicate the model moments are statistically indistinguishable from their empirical counterparts. A test of overidentifying restrictions delivers a p-value of 0.146, failing to reject the null hypothesis thereby lending additional support to the validity of the estimates.

5 Annual Smoothing and Overadjustments

Having evaluated the estimated model and assessed its fit to the targeted moments, we next turn to analyzing its ability to replicate the patterns of over- and underadjustments observed in the data.

5.1 Simulated Regression Coefficients

The model is able to successfully replicate the negative autocorrelation of revisions observed in the data. Figure 5 plots the autocorrelation of revisions across horizons both in the data and in the model. The simulated model-based estimates nearly always lie within 95% confidence interval of the estimated coefficients, both at the individual (top figure) and

Table 4: Model Estimation Results

Panel A: Parameter Estimates			
	Parameter	Estimate	Standard error
Persistence of latent state	ρ	0.535	0.028
State innovation dispersion	σ_w	2.369	0.210
Public signal noise	σ_e	1.657	0.154
Private signal noise	σ_v	0.721	0.139
Probability of quarterly update	q	0.929	0.065
Probability of annual update	p	0.609	0.052
Panel B: Moments			
	Model moment	Data moment	t-statistic
Standard deviation of nowcast	2.322	2.399	-0.602
Correlation of nowcast with annual forecast	0.747	0.745	-0.509
Standard deviation of annual forecast	1.577	1.622	-0.477
Standard deviation of revision	2.061	1.999	0.503
Correlation of revision with lagged error	0.139	0.147	-0.166
Standard deviation of lag error	2.101	2.128	-0.288
RMSE nowcast	2.129	2.150	-0.245
RMSE annual forecast	1.463	1.477	-0.225

Note: Panel A reports the model parameters with point estimates reported in the third column and standard errors reported in the fourth column. Panel B reports the model vs. data moments with t-statistics reported in the fourth column.

aggregate (bottom figure) levels.

Table 5 reports ten additional non-targeted moments. Panel A reports individual-level regression coefficients of errors-on-revisions at the current quarter as well as one- and two-quarter ahead horizons (rows 1-4). The fourth row of table reports the estimated coefficient from a regression of the annual (Q4/Q4) forecast error on the realized actual as in Kohlhas and Walther (2021). Across these four regressions, the model almost always predicts individual overreactions as in the top panel of Figure 5.[17]

Row 5 of Panel A reports estimates of forecast error persistence. We report this estimate to highlight our model's ability to reproduce another important feature of the data:

[17] The model does not generate a negative errors-on-revisions coefficient based on a regression of current-quarter errors on revisions. This is because the model assumes that the news forecasters receive is about today. As a result, forecasters place more importance on minimizing current quarter errors, and instead reshuffle their future forecasts to maintain annual consistency. If instead signals were informative about future quarters rather than the current quarter, then the model would generate a negative errors-on-revisions coefficient for current-quarter forecasts.

Figure 5: Model Fit to Revision Autocorrelation

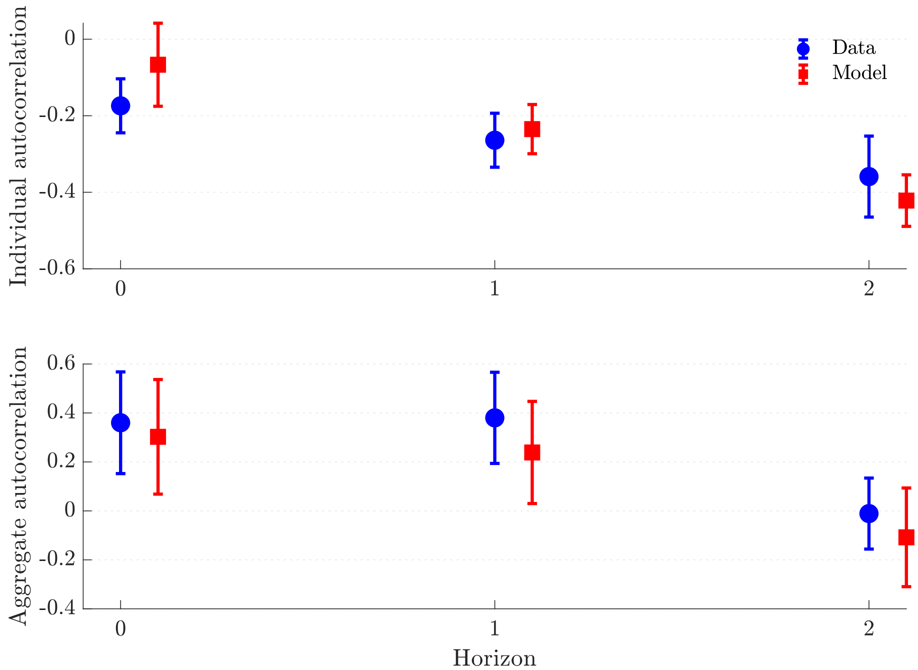

Note: The figure plots the empirical and model-implied autocorrelation of revisions. The top panel plots the individual autocorrelation coefficient, and the bottom panel plots the consensus-level autocorrelation coefficient. 95% confident intervals are displayed with each point estimate.

individual-level forecast error persistence. In a rational setting in which forecasters are able to observed lagged realizations of the variable of interest, errors should not exhibit persistence.[18] Our model is able to generate error persistence precisely because annual smoothing introduces lagged errors into reported forecasts. We find this to be a desirable feature of our model as it allows us to match this pattern in the data while allowing lagged realizations to reside the forecaster's information set.

Panel B of Table 5 report the aggregate analogs to the estimates in Panel A. The first three rows in particular, indicate that consensus forecasts exhibit inertia. Taken together, the results demonstrate that the model is successful in producing empirically relevant magnitudes of over- and underadjustments.

[18]The literature has often assumed that forecasters never observe the variable of interest, thereby preserving error persistence.

Table 5: Additional Non-targeted Moments

Panel A: Individual-Level	Model		Data	
$\beta(FECQ, FRCQ)$	0.047	(0.043)	-0.249	(0.071)
$\beta(FE1Q, FR1Q)$	-0.115	(0.099)	-0.228	(0.066)
$\beta(FE2Q, FR2Q)$	-0.431	(0.118)	-0.302	(0.051)
β(Annual FE, Realized actual)	-0.058	(0.080)	-0.089	(0.017)
$\beta(FECQ, FECQ_{-1})$	0.159	(0.046)	0.176	(0.046)
Panel B: Aggregate-Level	Model		Data	
$\beta(FECQ, FRCQ)$	0.378	(0.066)	0.354	(0.179)
$\beta(FE1Q, FR1Q)$	0.512	(0.220)	0.685	(0.313)
$\beta(FE2Q, FR2Q)$	0.347	(0.440)	0.667	(0.385)
β(Annual FE, Realized actual)	-0.058	(0.080)	-0.076	(0.064)
$\beta(FECQ, FECQ_{-1})$	0.086	(0.060)	0.083	(0.075)

Note: The table reports additional regression coefficients in the model as well in the data. Standard deviations and standard errors are reported in parentheses. 'FE' refers to forecast error, 'FR' refers to forecast revision, and 'CQ, 1Q, 2Q' refer to current quarter, one-quarter ahead, and two-quarters ahead, respectively.

5.2 Incorporating Non-Rational Expectations

To better understand the quantitative importance of our mechanism as a driver of overadjustments, we augment our model with a behavioral friction in a supplementary exercise. We choose a leading theory of non-rational expectations, diagnostic expectations (Bordalo et al., 2019; Bianchi et al., 2021a; Bordalo et al., 2021; Chodorow-Reich et al., 2021), which draws from the representativeness heuristic (Tversky and Kahneman, 1974). In particular, diagnostic forecasters place excess weight on new information such that their reported current-quarter prediction is:

$$x_{t|t}^{i,\theta} = \mathbb{E}_{it}(x_t) + \theta\big[\mathbb{E}_{it}(x_t) - \mathbb{E}_{it-1}(x_t)\big]$$

where θ is the degree of diagnosticity. When $\theta = 0$, the model collapses to the standard noisy information rational expectations model. On the other hand, in a world of diagnostic expectations, $\theta > 0$.

Figure 6: Annual Smoothing vs. Diagnostic Expectation Contributions

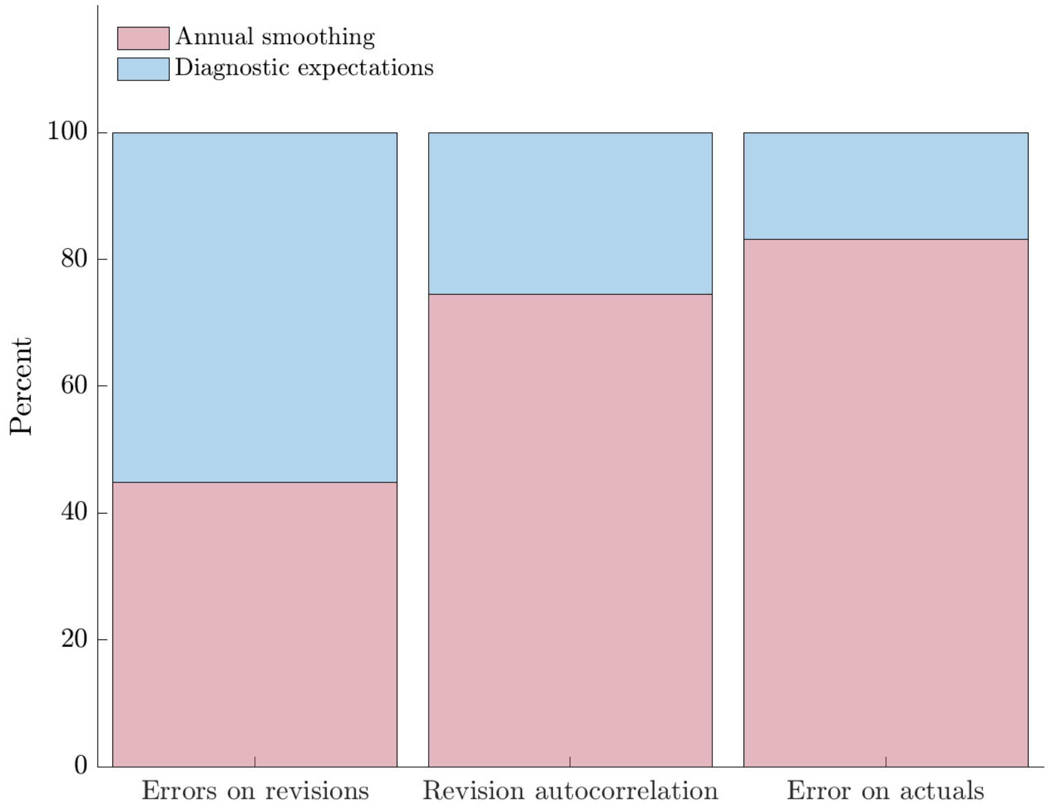

Note: The figure plots the contributions of annual smoothing and diagnostic expectations, in percent, to three measures of overreactions. The third measure is first studied in Kohlhas and Walther (2021).

The objective of this exercise is to model two channels of overreaction: (i) annual smoothing and (ii) diagnostic expectations, and to quantify the relative importance of our mechanism. To do so, we re-estimate the model with and without diagnostic expectations while targeting two additional moments: the contemporaneous covariance of errors and revisions, and the variance of contemporaneous errors. We add these moments to the estimation procedure in order to ensure that the model fits a well-known measure of overreactions, the coefficient of errors on revisions, as closely as possible.

The constrained model without diagnostic expectations, implies that the probability of

Case 2 updating, $q(1-p)$, is about 0.46.[19] When we incorporate diagnostic expectations, this probability falls to 0.31. At the same time, our estimate of diagnosticity is 0.45 which is 25% lower than the estimate obtained in Bordalo et al. (2020) using the same SPF data and a similar estimation approach. Overall, this indicates that annual smoothing and temporal consistency can contribute a meaningful amount of observed overreactions, even in the presence of other frictions.

We conclude the exercise by quantifying the importance of our mechanism through a decomposition exercise. Figure 6 displays three sets of stacked bars, each corresponding to a distinct measure of overadjustments. The red bar denotes the contribution of our annual smoothing motive to the overall measure of overadjustments while the blue bar denotes the contribution of diagnostic expectations. Once again, we find that annual smoothing is a meaningful, and in this case dominant, mechanism for generating overreactions. While there are a number of other plausible sources of overadjustments, these results suggest that annual smoothing can be a quantitatively important driver of overreactions.

5.3 Annual Smoothing Across SPF Variables

We next estimate our baseline model for the range of macroeconomic variables covered in the SPF. Table 6 reports empirical and simulated estimates of the revision autocorrelation, our non-targeted moment of choice. In general, we find that our model is broadly successful in reproducing the negatively autocorrelated revisions observed in the data.

5.4 Annual Smoothing By Forecaster Type

Annual inattention can arise due to reputational considerations or time and resource constraints associated with frequent model updating. While we cannot assess the reputational considerations hypothesis in the data, since the SPF is an anonymous survey, we can examine the plausibly of resource constraints. To do so, we exploit the SPF classifications of different

[19] Table D4 reports the parameter estimates.

Table 6: Estimates Across SPF Variables

	Revision Autocorrelation	
	Model	Data
Real GDP	**-0.230 (0.037)**	**-0.264 (0.036)**
Nominal GDP	**-0.169 (0.026)**	**-0.263 (0.032)**
Real Consumer Spending	**-0.325 (0.027)**	**-0.336 (0.089)**
GDP Deflator	**-0.165 (0.028)**	**-0.305 (0.052)**
Real residential investment	**-0.178 (0.027)**	**-0.230 (0.037)**
Real nonresidential investment	**-0.155 (0.027)**	**-0.186 (0.054)**
Real federal spending	**-0.203 (0.030)**	**-0.273 (0.061)**
Real state/local spending	**-0.178 (0.035)**	**-0.352 (0.071)**
Industrial production	**-0.209 (0.025)**	**-0.279 (0.035)**
CPI	**-0.447 (0.019)**	**-0.395 (0.088)**
Unemployment	-0.018 (0.057)	**-0.271 (0.048)**
Ten Year Bond	**-0.422 (0.019)**	**-0.396 (0.030)**
3-month bill	**-0.229 (0.045)**	**-0.301 (0.054)**
Housing starts	**-0.273 (0.026)**	**-0.331 (0.033)**

Note: The table reports one-quarter ahead revision autocorrelation coefficients in the model and the data for various macroeconomic variables covered in the SPF. Bold values are significantly negative at the 5% level.

forecaster types. Beginning in 1990, the SPF began collecting information on respondents' industries of employment. Respondents are labeled as either a financial service provider, a non-financial service provider, or neither. Financial service providers include asset managers, investment bankers, and insurance companies while non-financial forecasters include academics employed at universities, manufacturers, and consulting firms.[20] In general, a forecaster is able to switch across categories over time.

Assuming that the inattention is costlier among financial service providers, we hypothesize that these types of forecasters should exhibit less Case 2 forecasting from the lens of our model.[21] To assess this hypothesis, we re-estimate the baseline model for financial

[20]A full list is provided on page 33 of the SPF documentation: https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/spf-documentation.pdf?la=en&hash=F2D73A2CE0C3EA90E71A363719588D205

[21]Financial service providers could face plausibly higher inattention costs relative to non-financial service providers for a number of reasons. Many of the client demands among financial service providers require operating with up-to-date information. On the other hand, academics or other non-financial service providers must complete other tasks on a regular basis that might not require highly updated information on macroe-

Figure 7: Model Fit to Revision Autocorrelation, by Forecaster Type

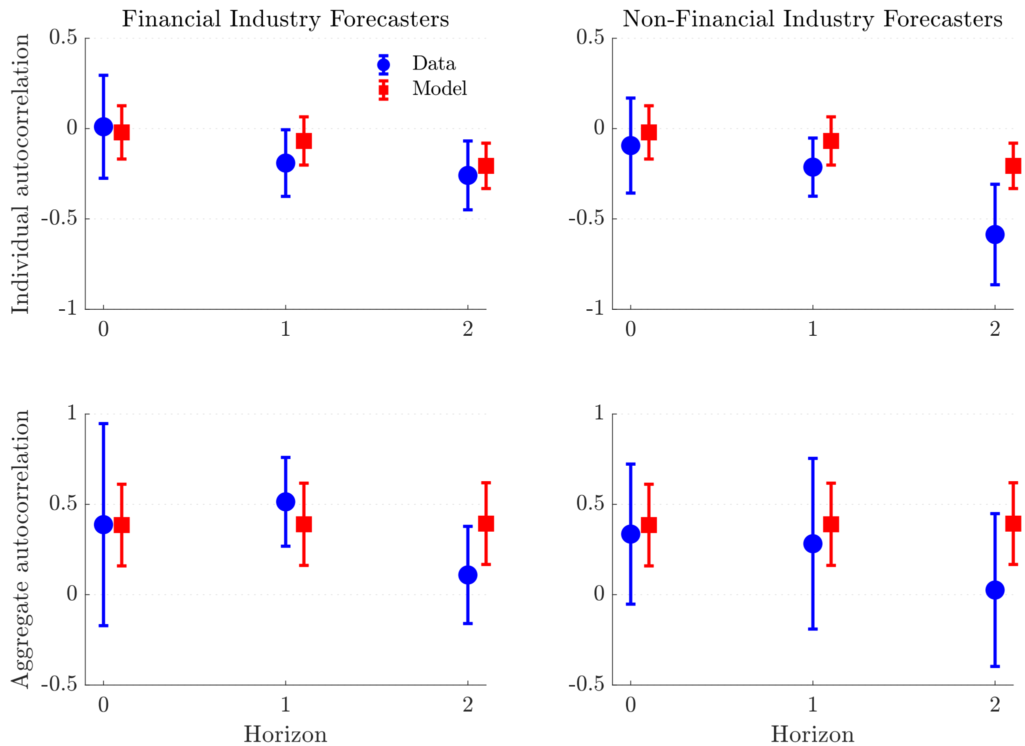

Note: The figure plots the empirical and model-implied autocorrelation of revisions. The top two panels plot the individual autocorrelation coefficient, and the bottom two panels plot the consensus-level autocorrelation coefficient. The first column of panels refers to the financial industry forecaster subsample while the second column of panels refers to non-financial forecaters. 95% confident intervals are displayed with each point estimate.

and non-financial forecasters separately and find the strongest evidence of annual smoothing among non-financial service providers.[22] Figure 7 displays the non-targeted fit of the estimated models to the autocorrelation of revisions. Overall, our estimated models are able to successfully match the autocorrelation of revisions across either type of forecaster, which are more strongly negative among non-financial forecasters. These results favor a time or resource-based interpretation of annual inattention.

conomic developments.

[22]The estimation results are reported in Table D5 in Appendix D.

Table 7: Information Frictions Across Models

	Probability of updating	Implied friction	Sticky info contribution	Noisy info contribution
Panel A: Real GDP				
Quarterly	0.929	0.143	50.1%	49.9%
Annual	0.609	0.438	89.4%	10.6%
Panel B: Inflation				
Quarterly	0.971	0.175	16.8%	83.2%
Annual	0.432	0.633	89.8%	10.2%

Note: The table reports estimated updating probabilities, implied information frictions, and contributions of sticky and noisy information for real GDP and inflation at quarterly and annual frequencies. Implied information frictions are computed based on (6) with model-implied Kalman gains $\{0.917, 0.007\}$ and $\{0.800, 0.051\}$ for real GDP and inflation, respectively. Contributions of sticky and noisy information are computed according to (7)

6 Implications for Information Frictions

In addition to serving as a source of observed overadjustments, our model can also speak to the literature on information frictions. Since our model does not allow us to readily extract a coefficient of information rigidity from an OLS regression, we simulate the estimated model in order to quantify the size of information frictions.

6.1 Model-Implied Information Rigidities

Column 3 of Table 7 reports measures of implied information rigidity for SPF forecasts of real GDP and inflation. Since our model is a hybrid sticky-noisy information model, we define the implied information friction to be:

$$\text{Implied friction} = \big[1 - \Pr(\text{update})\big] + \Pr(\text{update}) \times (1 - \kappa_1 - \kappa_2) \tag{6}$$

where Pr(update) denotes the probability of updating, which reflects the sticky information feature of the model. This probability varies across quarterly and annual frequencies. Moreover, the role of noisy information in overall information frictions is understood through the coefficients $\{\kappa_1, \kappa_2\}$ which denote the Kalman gains.[23] In traditional models of either sticky information or noisy information, the relevant information rigidity is governed by either the probability of updating or the Kalman gain(s). Here, the implied friction is a combination of these two objects. With some probability, forecasters do not update. In this case, they effectively place a weight of zero on new information ($\kappa_1 = \kappa_2 = 0$) so that the first term in (6) is scaled by one. With some probability, forecasters do update, in which case they weigh new information based on the Kalman gains. Upon updating, the relevant information friction is one minus the value of these optimal weights. Together, these terms capture the general notion of an information friction in a hybrid sticky-noisy information model.

To compare our implied information frictions to estimates documented in the literature, we also report model estimates using inflation forecasts based on the GDP deflator. At a quarterly frequency, we estimate information frictions to be about 0.18 while, for annual forecasts, we find that information frictions are higher, at 0.63. For reference, Coibion and Gorodnichenko (2015) estimate coefficients of information rigidity to be around 0.54 while to Ryngaert (2017) estimates information frictions to be roughly 0.33. Importantly, whereas existing estimates imply a single information friction for all frequencies, our analysis indicates that there is a difference in information rigidities across quarterly and annual frequencies. We note that the average of our quarterly and annual implied information frictions hover around these previously documented estimates.

6.2 Contributions of Sticky and Noisy Information

The literature on survey expectations has documented evidence consistent with both sticky and noisy information. Our results indicate that the data favor a hybrid model featuring signal extraction and frequency-specific sticky information. In addition to providing estimates of information frictions based on both sticky and noisy information, our model can also quantify the relative importance of each of these channels. To do so, we normalize the

[23] In particular, κ_1 denotes the weight placed on the private contemporaneous signal and κ_2 is the weight placed on the lagged realization of the macro variable.

implied information friction to equal one

$$1 = \underbrace{\frac{1 - \Pr(\text{update})}{[1 - \Pr(\text{update})] + \Pr(\text{update}) \times (1 - \kappa_1 - \kappa_2)}}_{\text{Sticky info contribution}} + \underbrace{\frac{\Pr(\text{update}) \times (1 - \kappa_1 - \kappa_2)}{[1 - \Pr(\text{update})] + \Pr(\text{update}) \times (1 - \kappa_1 - \kappa_2)}}_{\text{Noisy info contribution}}. \tag{7}$$

The first term in the above expression quantifies the role of sticky information in the overall measured information rigidity while the second term quantifies the importance of signal processing. The final two columns of Table 7 report the contributions of each forms of imperfect information to the implied friction reported in column 3. The results from this accounting exercise suggest that noisy information is the primary contributor to estimated information frictions among quarterly inflation forecasts, while sticky information becomes substantially more important at the annual frequency.

7 Conclusion

We show that professional forecasters exhibit over- and underadjustments, and they appear to offset their updates. We build a hybrid sticky-noisy information model that can account for these facts. From the lens of our model, overreactions arise because of annual smoothing and temporal consistency. When faced with new information, forecasters offset their current updates further along their annual trajectories. The tradeoff between minimizing errors and satisfying consistency generates a negative autocorrelation of revisions. The estimated model successfully fits key micro moments among professional forecasters, and can explain a meaningful amount of overreactions to real GDP as well as other variables in the SPF. When comparing our model to a theory of non-rational explanations, we find that annual smoothing explains much of the overreaction among professional forecasters.

Our results also imply that information frictions and their composition vary by frequency. Our model mainly attributes the higher information rigidity at the annual level to information stickiness, while the smaller quarterly frictions are mainly attributed to noisy information. Future research might be able to provide a deeper microfoundation for annual smoothing, be it rational or behavioral. For policymakers, the result that quarterly predictions are updated almost every quarter, and are contaminated by overadjustments due to annual smoothing, implies that they should focus on managing medium- and long-term expectations as they are more informative and persistent.

References

Afrouzi, H., Kwon, S., Landier, A., Ma, Y., and Thesmar, D. (2021). Overreaction in Expectations: Evidence and Theory. Working Paper.

Andrade, P. and Bihan, H. L. (2013). Inattentive Professional Forecasters. *Journal of Monetary Economics*, 60:967–982.

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the Sensitivity of Parameter Estimates to Estimation Moments. *The Quarterly Journal of Economics*, 132:1553–1592.

Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2020). Imperfect Macroeconomic Expectations; Evidence and Theory. *NBER Macroeconomics Annual 2020*, 35.

Azeredo da Silvera, R., Sung, Y., and Woodford, M. (2020). Optimally Imprecise Memory and Biased Forecasts. NBER Working Paper 28075.

Baker, S. R., McElroy, T. S., and Sheng, X. S. (2020). Expectation formation following large, unexpected shocks. *Review of Economics and Statistics*, 102(2):287–303.

Bianchi, F., Ilut, C., and Saijo, H. (2021a). Diagnostic Business Cycles. NBER Working Paper 28604.

Bianchi, F., Ludvigson, S. C., and Ma, S. (2021b). Belief Distortions and Macroeconomic Fluctuations. NBER Working Paper 27406.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in Macroeconomic Expectations. *American Economic Review*, 110:2748–2782.

Bordalo, P., Gennaioli, N., Porta, R. L., and Shleifer, A. (2019). Diagnostic expectations and stock returns. *The Journal of Finance*, 74(6):2839–2874.

Bordalo, P., Gennaioli, N., Schleifer, A., and Terry, S. (2021). Real Credit Cycles. NBER Working Paper 28416.

Broer, T. and Kohlhas, A. (2019). Forecaster (Mis-)Behavior. Working Paper.

Bürgi, C. R. S. (2016). What Do We Lose When We Average Expectations? Working Papers 2016-013, The George Washington University, Department of Economics, H. O. Stekler Research Program on Forecasting.

Bürgi, C. R. S. (2017). Bias, rationality and asymmetric loss functions. *Economics Letters*, 154:113–116.

Bürgi, C. R. S. and Ortiz, J. L. (2021). Temporal consistency of forecasts and data releases. *Mimeo*.

Chodorow-Reich, G., Guren, A. M., and McQuade, T. J. (2021). The 2000s Housing Cycle with 2020 Hindsight: A Neo-Kindlebergerian View. NBER Working Paper 29140.

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

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 109:465–490.

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

Dovern, J., Fritsche, U., Loungani, P., and Tamirisa, N. (2015). Information Rigidities: Comparing Average and Individual Forecasts for a Large International Panel. *International Journal of Forecasting*, 31:144–154.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data. *The Review of Economics and Statistics*, 80:549–560.

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

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

Kucinskas, S. and Peters, F. (2019). Measuring Biases in Expectation Formation. Working Paper.

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

Nordhaus, W. D. (1987). Forecasting Efficiency: Concepts and Applications. *The Review of Economics and Statistics*, 69:667–674.

Ryngaert, J. (2017). What Do (and Don't) Forecasters Know About U.S. Inflation? Working Paper.

Scotese, C. A. (1994). Forecast Smoothing and the Optimal Under-utilization of Information at the Federal Reserve. *Journal of Macroeconomics*, 16:653–670.

Sims, C. (2003). The Implications of Rational Inattention. *Journal of Monetary Economics*, 50:665–690.

Tversky, A. and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185:1124–1131.

Woodford, M. (2001). Imperfect Common Knowledge and the Effects of Monetary Policy. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps*.

Appendix A Empirics

A.1 Summary Statistics

We utilize data from the SPF spanning 1968Q4-2019Q4. Table A1 report summary statistics of real GDP forecast errors and revisions across horizons, as well as real-time outcomes and data revisions.

Table A1: SPF Real GDP Summary Statistics

	Mean	Median	Std. deviation	25%	75%
Forecasts					
Current quarter	2.364	2.524	2.786	1.416	3.603
One quarter ahead	2.679	2.700	2.437	1.900	3.615
Two quarters ahead	2.906	2.836	2.220	2.082	3.776
Forecast errors					
Current quarter	0.021	-0.037	2.460	-1.261	1.260
One quarter ahead	-0.310	-0.254	2.923	-1.664	1.068
Two quarters ahead	-0.673	-0.348	4.108	-1.826	1.091
Forecast revisions					
Current quarter	-0.269	-0.119	2.234	-0.994	0.565
One quarter ahead	-0.166	-0.036	2.037	-0.637	0.407
Two quarters ahead	-0.170	-0.030	2.035	-0.558	0.325
Real GDP					
Real-time outcomes	2.408	2.529	2.969	1.310	3.885
Data revisions	-0.023	-0.072	0.760	-0.523	0.332

Note: The table reports summary statistics for the relevant variables utilized in the main text. The sample is constructed from SPF real GDP growth forecast data. The unbalanced panel spans 1968Q4-2019Q4, and consists of 253 unique forecasters.

A.2 Additional Evidence of Over- and Underreaction

In addition to the regression results presented Section2, Kohlhas and Walther (2021) provide an additional measure of overreaction based on regressing ex-post forecast errors on outcomes. Table A2 reports the results from this regression based on our sample.

Table A2: Overreaction to Realized Output

	Annual error	Annual error
Realized outcome	-0.085***	-0.089***
	(0.016)	(0.017)
Fixed effects	None	Forecaster
Observations	4559	4540

Note: The table reports panel regression results from SPF forecasts of real GDP based on the regression of errors on realized output in Kohlhas and Walther (2021). *** denotes 1% significance, ** denotes 5% significance, and * denotes 10% significance.

Appendix B Model

B.1 Deriving the Reported Forecast

Suppose that in each period, professional forecasters devise predictions across a number of horizons, H. Forecasters in the model wish to minimize the sum of their mean square errors:

$$\min_{\{\widehat{x}^i_{t+h|t|t+k}\}} \sum_{h=0}^{H} (x_{t+h} - \widehat{x}^i_{t+h|t+k})^2, \tag{8}$$

where $\widehat{x}^i_{t+h|t+k}$ denote forecaster i's predictions about x_t h-steps into the future, based on information at time $t+k$.

When forecasters are able to freely update quarterly and annual forecasts, they report

$$\widehat{x}^i_{t+k'|t+k} = \mathbb{E}_{it+k}(x_{t+k'}) \;\; \forall k' \in [0, H], \quad \text{and} \quad \frac{1}{H+1}\sum_{h=0}^{H} \widehat{x}^i_{t+h|t+k}$$

as their quarterly and annual forecasts, respectively.

If the forecaster is able to update her short-run predictions but not her long-run predictions, then she must solve the optimization problem above subject to the requirement that the updated quarterly forecasts coincide with the outdated annual forecast:

$$\frac{1}{H+1}\sum_{h=0}^{H} \widehat{x}^i_{t+h|t+k} = \frac{1}{H+1}\sum_{h=0}^{H} \widehat{x}^i_{t+h|t+k-j}, \tag{9}$$

where j denotes the period in which the annual forecast was last updated. In this case, the forecaster solves (8) subject to (9).

The Lagrangian is

$$\mathcal{L} = \sum_{h=0}^{H} (x_{t+h} - \widehat{x}^i_{t+h|t+k})^2 - \lambda\left(\frac{1}{H+1}\sum_{h=0}^{H} \widehat{x}^i_{t+h|t+k} - \frac{1}{H+1}\sum_{h=0}^{H} \widehat{x}^i_{t+h|t+k-j}\right)$$

The first order condition with respect to the reported forecast $\widehat{x}^i_{t+k'|t+k}$ implies

$$\widehat{x}^i_{t+k'|t+k} = \mathbb{E}_{it+k}(x_{t+k'}) + \frac{\lambda}{2(H+1)}. \tag{10}$$

Combining the FOC with the definition of the constraint delivers:

$$\frac{1}{H+1}\sum_{h=0}^{H}\widehat{x}^i_{t+h|t+k-j} = \frac{1}{H+1}\sum_{h=0}^{H}\left[\mathbb{E}_{it+h}(x_{t+k'}) + \frac{\lambda}{2(H+1)}\right].$$

Rearranging, we obtain:

$$\lambda = 2(H+1)\left[\frac{1}{H+1}\sum_{h=0}^{H}\widehat{x}^i_{t+h|t+k-j} - \frac{1}{H+1}\sum_{h=0}^{H}\mathbb{E}_{it+k}(x_{t+k'})\right]$$

Substituting this expression for the Lagrange multiplier into the FOC for the reported forecast, we obtain an intuitive expression:

$$\widehat{x}^i_{t+k'|t+k} = \mathbb{E}_{it+k}(x_{t+k'}) + \left[\frac{1}{H+1}\sum_{h=0}^{H}\widehat{x}^i_{t+k'|t+k-j} - \frac{1}{H+1}\sum_{h=0}^{H}\mathbb{E}_{it+k}(x_{t+k'})\right]$$

or, equivalently,[24]

$$\widehat{x}^i_{t+k'|t+k} = \mathbb{E}_{it+k}(x_{t+k'}) + \frac{1}{H+1}\sum_{h=0}^{H}\left[\mathbb{E}_{it+k-j}(x_{t+k'}) - \mathbb{E}_{it+k}(x_{t+k'})\right]. \tag{11}$$

B.2 Analytical Characterization of Individual Revisions

We can split the second term in (11) to distinguish events that have come to pass (before $t+k$) from events that have yet to pass (on or after $t+k$):

[24] This follows from the fact that whenever the forecaster constructed her outdated annual, she did so optimally, based on the conditional expectation as of date $t+k-j$.

$$
\begin{aligned}
\widehat{x}^i_{t+k'|t+k} = \mathbb{E}_{it+k}(x_{t+k'}) &+ \frac{1}{H+1}\sum_{h=0}^{k-1}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - x_{t+h}\right] \\
&+ \frac{1}{H+1}\sum_{h=k}^{H}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - \mathbb{E}_{it+k}(x_{t+h})\right].
\end{aligned}
$$

The middle term reflects lagged forecast errors while the last term reflects forecast revisions. The previous period forecast for the same horizon $(t+k')$ is

$$
\begin{aligned}
\widehat{x}^i_{t+k'|t+k-1} = \mathbb{E}_{it+k-1}(x_{t+k'}) &+ \frac{1}{H+1}\sum_{h=0}^{k-2}\left[\mathbb{E}_{it+k-\ell}(x_{t+h}) - x_{t+h}\right] \\
&+ \sum_{h=k-1}^{H}\left[\mathbb{E}_{it+k-\ell}(x_{t+h}) - \mathbb{E}_{it+k-1}(x_{t+h})\right].
\end{aligned}
$$

The forecast revision is therefore:

$$
\begin{aligned}
\widehat{x}^i_{t+k'|t+k} - \widehat{x}^i_{t+k'|t+k-1} = {} & \mathbb{E}_{it+k}(x_{t+k'}) - \mathbb{E}_{it+k-1}(x_{t+k'}) \\
&+ \frac{1}{H+1}\left\{\sum_{h=0}^{k-1}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - x_{t+h}\right] - \sum_{h=0}^{k-2}\left[\mathbb{E}_{it+k-\ell}(x_{t+h}) - x_{t+h}\right]\right\} \\
&+ \frac{1}{H+1}\left\{\sum_{h=k}^{H}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - \mathbb{E}_{it+k}(x_{t+h})\right] - \sum_{h=k-1}^{H}\left[\mathbb{E}_{it+k-\ell}(x_{t+h}) - \mathbb{E}_{it+k-1}(x_{t+h})\right]\right\},
\end{aligned}
$$

Or, equivalently,

$$
\begin{aligned}
\widehat{x}^i_{t+k'|t+k} - \widehat{x}^i_{t+k'|t+k-1} = {} & \mathbb{E}_{it+k}(x_{t+k'}) - \mathbb{E}_{it+k-1}(x_{t+k'}) \\
&- \frac{1}{H+1}\left\{\left[x_{t+k-1} - \mathbb{E}_{it+k-j}(x_{t+k-1})\right] + \sum_{h=0}^{k-2}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - \mathbb{E}_{it+k-\ell}(x_{t+h})\right]\right\} \\
&+ \frac{1}{H+1}\left\{\left[\mathbb{E}_{it+k-1}(x_{t+k-1}) - \mathbb{E}_{it+k-\ell}(x_{t+k-1})\right] + \sum_{h=k}^{H}\left[\mathbb{E}_{it+k}(x_{t+h}) - \mathbb{E}_{it+k-1}(x_{t+h})\right]\right. \\
&\left. - \sum_{h=k}^{H}\left[\mathbb{E}_{it+k-j}(x_{t+h}) - \mathbb{E}_{it+k-\ell}(x_{t+h})\right]\right\}
\end{aligned}
$$

B.3 Initial Underreaction, Delayed Overreaction

The model is also capable of generating the initial underreaction and delayed overreaction dynamics recently documented in the literature (Angeletos et al., 2020; Bianchi et al., 2021b). Figure B1 plots the simulated response of aggregate forecast errors to shock to the underlying state, s_t, in our model (solid red line) and a model without annual inattention (dashed blue line). The figure corroborates the intuition highlighted above by demonstrating that errors inherit the excess volatility in revisions that arise with Case 2 updating.

Figure B1: Dynamics of Annual Smoothing

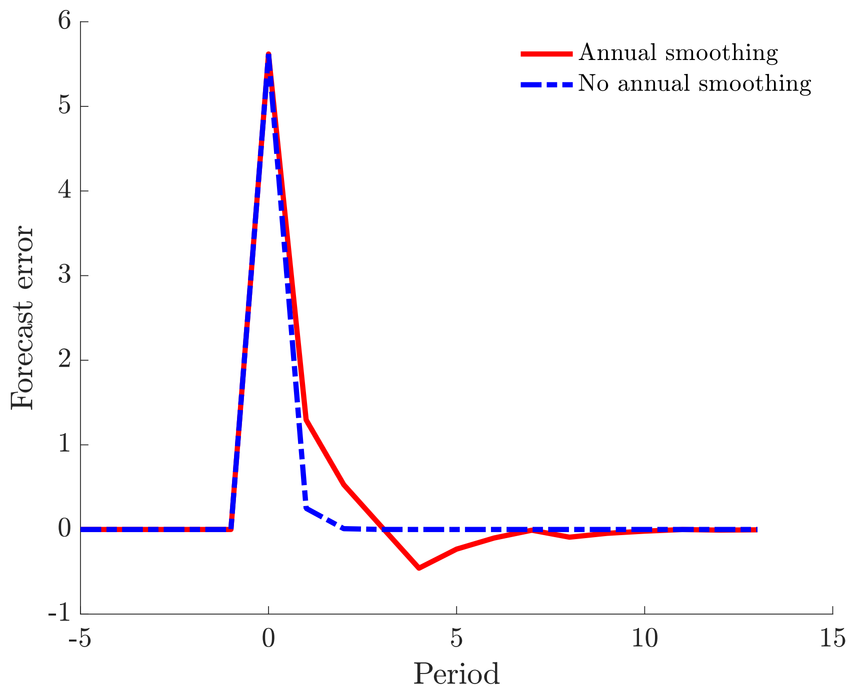

Note: The figure plots simulated impulse responses to a one standard deviation shock to w_t in the model. The solid red line plots the impulse response in our model with annual smoothing and temporal consistency while the dashed blue line plots the impulse response without annual smoothing.

Appendix C Estimation

The model is estimated via the simulated method of moments. Operationally, this is done by simulating a balanced panel of 250 forecasters over 40 periods, consistent with the average number of quarterly forecasts that a unique forecaster contributes throughout the history of the survey.[25] For each iteration, the target moments are computed, averaged across simulations, and compared to their empirical analogs. The six-dimensional parameter vector, θ, is selected to minimize the weighted distance between simulated moments and empirical moments, where the asymptotically efficient weighting matrix is specified.

Formally, we search the parameter space via a particle swarm routine to find the $\widehat{\theta}$ that minimizes the following objective

$$\min_{\theta} \big(m(\theta) - m(X)\big)' W \big(m(\theta) - m(X)\big)$$

where $m(\theta)$ denotes the simulated moments, $m(X)$ denotes the empirical moments, and W denotes the weighting matrix. The limiting distribution of the estimated parameter vector $\widehat{\theta}$ is

$$\sqrt{N}(\widehat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, \Sigma)$$

where

$$\Sigma = \left(1 + \frac{1}{S}\right)\left[\left(\frac{\partial m(\theta)}{\partial \theta}\right)' W \left(\frac{\partial m(\theta)}{\partial \theta}\right)\right]^{-1}$$

and $S = 100$. Standard errors are obtained by numerically computing the partial derivative of the simulated moment vector with respect to the parameter vector.

C.1 Identification

The eight moments jointly determine the six parameters that reside in vector θ. To supplement the discussion on monotone relationships reported in Figure C2, we additionally

[25] Similar results are obtained when mimicking the unbalanced nature of the panel data by simulating a larger set of forecasters and matching missing observations.

report the sensitivity of each of the parameters to changes in each of the moments in Figure C3. These sensitivities are an implementation of Andrews et al. (2017). In particular, the sensitivity of $\widehat{\theta}$ to $m(\theta)$ is

$$\Lambda = -\left[\left(\frac{\partial m(\theta)}{\partial\theta}\right)' W \left(\frac{\partial m(\theta)}{\partial\theta}\right)\right]^{-1} \left(\frac{\partial m(\theta)}{\partial\theta}\right)' W.$$

I transform this matrix so that the estimates can be interpreted as elasticities of the parameters with respect to moments.

Figure C2: Comparative Statics

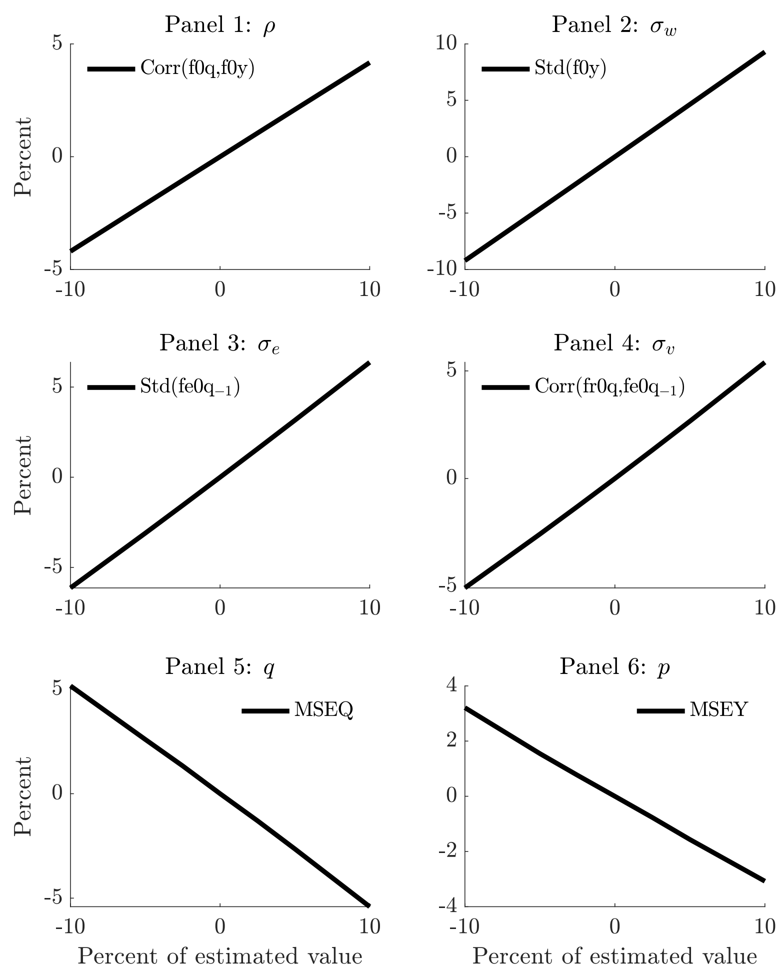

Note: Each panel displays a monotonic relationship between the parameter on the horizontal axis and a given moment (or pair of moments). The vertical axis measures the percent deviation of the given moment from its estimated value in Table 4.

Figure C3: Sensitivity

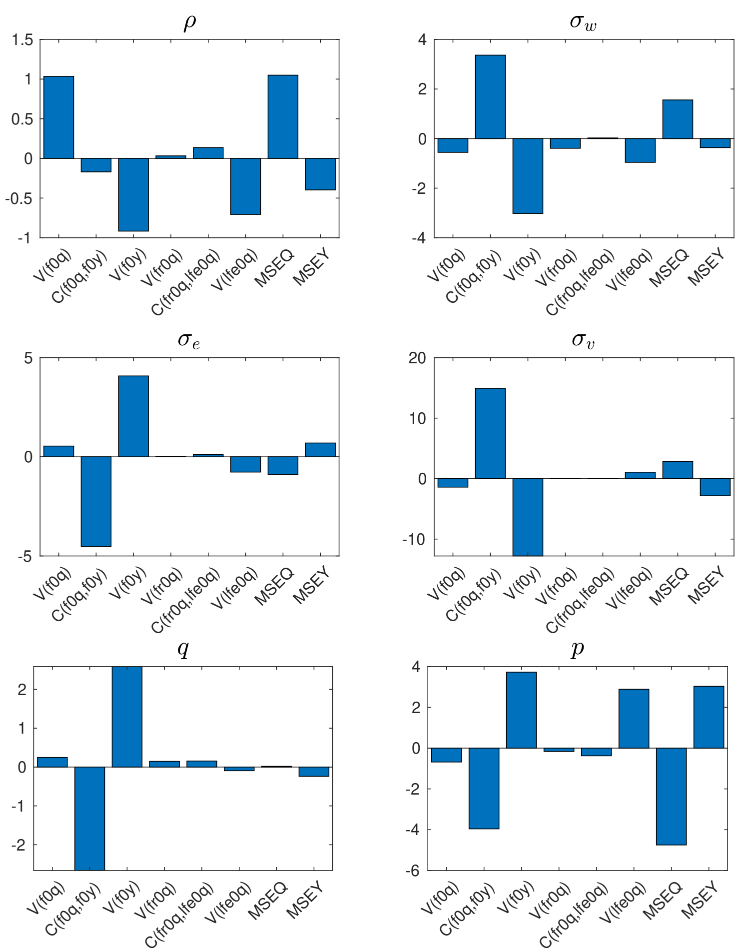

Note: The figure computes the elasticities of estimated parameters to moments as in Andrews et al. (2017).

Appendix D Additional Robustness Exercises

In this section, we consider a variety of robustness checks. First, we examine the role that rounding plays in the parameter estimates. We then augment our model with diagnostic expectations to assess the relative importance of our mechanism in generating overadjustments. Next, we report the results for across forecaster type, and then report the estimates based on real GDP forecasts from the Bloomberg Survey as well as SPF inflation forecasts. Following this, we undertake a sub-sample analysis, estimating the baseline model before and after 1990. Finally, we consider an alternate data generating process for the underlying state.

D.1 Rounding

We first report parameter estimates under the assumption that forecasters round their predictions to the nearest 0.10 percentage point. We find that this rounding assumption does not meaningfully change our parameter estimates (see Table D3).[26]

[26] Studying more traditional Gaussian measurement error introduces an identification problem between the measurement error dispersion and private signal noise dispersion, σ_v. At the same time, rounding is a well understood phenomenon in survey expectations. For this reason, we focus on this form of measurement error.

Table D3: Model Estimation Results, Rounding Reported Forecasts to Nearest 0.1pp

Panel A: Parameter Estimates			
	Parameter	Estimate	Standard error
Persistence of latent state	ρ	0.534	0.028
State innovation dispersion	σ_w	2.372	0.209
Public signal noise	σ_e	1.659	0.154
Private signal noise	σ_v	0.725	0.139
Probability of quarterly update	q	0.929	0.065
Probability of annual update	p	0.609	0.052
Panel B: Moments			
	Model moment	Data moment	t-statistic
Standard deviation of nowcast	2.322	2.399	-0.599
Correlation of nowcast with annual forecast	0.747	0.745	-0.518
Standard deviation of annual forecast	1.576	1.622	-0.483
Standard deviation of revision	2.063	1.999	0.519
Correlation of revision with lagged error	0.138	0.147	-0.162
Standard deviation of lag error	2.104	2.128	-0.260
RMSE nowcast	2.131	2.150	-0.214
RMSE annual forecast	1.464	1.477	-0.210

Note: Panel A reports the model vs. data moments with t-statistics reported in the fourth column. Panel B reports the model parameters with point estimates reported in the third column and standard errors reported in the fourth column.

D.2 Incorporating Diagnostic Expectations

Table D4: Model Estimation Results, Diagnostic Expectations

	Parameter	Unconstrained	Constrained
Persistence of latent state	ρ	0.606	0.634
		(0.034)	(0.035)
State innovation dispersion	σ_w	1.935	1.853
		(0.192)	(0.184)
Public signal noise	σ_e	1.324	1.289
		(0.167)	(0.183)
Private signal noise	σ_v	0.041	1.011
		(0.008)	(0.175)
Probability of quarterly update	q	0.655	1.000
		(0.083)	(0.116)
Probability of annual update	p	0.528	0.537
		(0.059)	(0.066)
Diagnosticity	θ	0.451	0.000
		(0.083)	-

Note: The table reports parameter estimates of the baseline model with and without diagnostic expectations. The "Unconstrained" column refers to the full model with annual inattention and diagnostic expectations. The "Constrained" column refers to the model with only annual inattention. Standard errors are reported in parentheses.

D.3 Financial vs. Non-Financial Forecasters

Table D5: Model Estimation Results, By Forecaster Type

	Parameter	Financial Industry	Non-financial Industry
Persistence of latent state	ρ	0.491	0.617
		(0.114)	(0.060)
State innovation dispersion	σ_w	1.365	1.339
		(0.137)	(0.120)
Public signal noise	σ_e	1.162	0.958
		(0.240)	(0.239)
Private signal noise	σ_v	1.015	1.130
		(0.142)	(0.123)
Probability of quarterly update	q	0.999	0.999
		(0.113)	(0.527)
Probability of annual update	p	0.989	0.435
		(0.432)	(0.173)

Note: The table reports parameter estimates of the baseline model, estimated separately over a sample of financial industry forecasters and non-financial industry forecasters, respectively.

D.4 Bloomberg Real GDP Forecasts

Table D6: Model Estimation Results, Real GDP Forecasts

Panel A: Parameter Estimates			
	Parameter	Estimate	Standard error
Persistence of latent state	ρ	0.613	0.042
State innovation dispersion	σ_w	1.169	0.060
Public signal noise	σ_e	1.045	0.046
Private signal noise	σ_v	0.002	0.0003
Probability of quarterly update	q	0.888	0.049
Probability of annual update	p	0.736	0.077
Panel B: Moments			
	Model moment	Data moment	t-statistic
Standard deviation of nowcast	1.322	1.759	-6.427
Correlation of nowcast with annual forecast	0.794	0.743	-4.835
Standard deviation of annual forecast	0.979	1.213	-4.603
Standard deviation of revision	1.099	1.222	-2.839
Correlation of revision with lagged error	0.099	0.160	-2.648
Standard deviation of lag error	1.178	1.195	-1.048
RMSE nowcast	1.191	1.198	-0.414
RMSE annual forecast	0.763	0.827	-2.423

Note: Panel A reports the model vs. data moments with t-statistics reported in the fourth column. Panel B reports the model parameters with point estimates reported in the third column and standard errors reported in the fourth column.

D.5 Inflation Forecasts

Table D7: Model Estimation Results, Inflation Forecasts (Deflator)

Panel A: Parameter Estimates			
	Parameter	Estimate	Standard error
Persistence of latent state	ρ	0.638	0.044
State innovation dispersion	σ_w	1.546	0.114
Public signal noise	σ_e	0.877	0.152
Private signal noise	σ_v	0.793	0.127
Probability of quarterly update	q	0.971	0.079
Probability of annual update	p	0.432	0.056
Panel B: Moments			
	Model moment	Data moment	t-statistic
Standard deviation of nowcast	1.559	1.603	-0.510
Correlation of nowcast with annual forecast	0.796	0.777	-0.315
Standard deviation of annual forecast	1.135	1.176	-0.558
Standard deviation of revision	1.258	1.407	-1.588
Correlation of revision with lagged error	0.200	0.225	-0.964
Standard deviation of lag error	1.405	1.505	-1.467
RMSE nowcast	2.048	2.294	-1.349
RMSE annual forecast	1.165	1.371	-1.343

Note: Panel A reports the model vs. data moments with t-statistics reported in the fourth column. Panel B reports the model parameters with point estimates reported in the third column and standard errors reported in the fourth column.

D.6 Sub-sample Analysis (Pre- and Post-1990)

The SPF, as well as broader macroeconomic dynamics, experienced important changes between 1968-2019. In this section, we estimate the model for two sub-periods: 1968-1989 (Table D8) and 1990-2019 (Table D9). Overall, we find that our headline conclusions hold across the sub-samples with the estimated parameters differing across samples as expected. For instance, we estimate underlying state to be less persistent and more volatile between 1969-2019.

Table D8: Model Estimation Results (1968-1989)

Panel A: Parameter Estimates			
	Parameter	Estimate	Standard error
Persistence of latent state	ρ	0.498	0.034
State innovation dispersion	σ_w	3.684	0.499
Public signal noise	σ_e	2.493	0.495
Private signal noise	σ_v	0.566	0.128
Probability of quarterly update	q	0.915	0.122
Probability of annual update	p	0.615	0.064
Panel B: Moments			
	Model moment	Data moment	t-statistic
Standard deviation of nowcast	3.552	3.678	-0.693
Correlation of nowcast with annual forecast	0.728	0.746	-0.919
Standard deviation of annual forecast	2.366	2.483	-0.773
Standard deviation of revision	3.189	3.101	0.427
Correlation of revision with lagged error	0.098	0.110	-0.156
Standard deviation of lag error	3.084	3.050	0.220
RMSE nowcast	3.127	3.215	-0.554
RMSE annual forecast	2.169	2.299	-0.976

Note: Panel A reports the model vs. data moments with t-statistics reported in the fourth column. Panel B reports the model parameters with point estimates reported in the third column and standard errors reported in the fourth column.

Table D9: Model Estimation Results (1990-2019)

Panel A: Parameter Estimates			
	Parameter	Estimate	Standard error
Persistence of latent state	ρ	0.587	0.022
State innovation dispersion	σ_w	1.400	0.117
Public signal noise	σ_e	1.044	0.087
Private signal noise	σ_v	1.009	0.181
Probability of quarterly update	q	0.998	0.065
Probability of annual update	p	0.478	0.055
Panel B: Moments			
	Model moment	Data moment	t-statistic
Standard deviation of nowcast	1.308	1.472	-3.037
Correlation of nowcast with annual forecast	0.766	0.741	-1.460
Standard deviation of annual forecast	0.957	0.992	-0.796
Standard deviation of revision	1.115	1.181	-1.602
Correlation of revision with lagged error	0.195	0.243	-1.879
Standard deviation of lag error	1.497	1.520	-0.647
RMSE nowcast	1.516	1.551	-0.992
RMSE annual forecast	0.969	0.994	-0.821

Note: Panel A reports the model vs. data moments with t-statistics reported in the fourth column. Panel B reports the model parameters with point estimates reported in the third column and standard errors reported in the fourth column.

D.7 Alternate Data Generating Process

Whereas offsetting revisions can be an artifact of annual smoothing, these patterns could also arise under a more general data generating process. If so, then we might be erroneously attributing the empirical finding to annual smoothing. In this section, we provide results in support of our mechanism under richer dynamics.

We extend our model to feature an AR(2) process for real GDP growth. We select an AR(2) process for three reasons. First, we find that the AR(2) fits real GDP growth best in the sense that it delivers the lowest information criteria.[27] Second, an AR(2) is highly feasible to estimate with the current SMM approach as it only adds one parameter to the model. Third, an AR(2) allows us to remain consistent with others in the literature who similarly examine alternate data generating processes for their models (Bordalo et al., 2020).

The key modification relative to the baseline model detailed in the main text is that the underlying latent state now evolves as follows:

$$s_t = (1 - \rho_1 - \rho_2)\mu + \rho_1 s_{t-1} + \rho_2 s_{t-2} + w_t, \quad w_t \sim N(0, \sigma_w^2)$$

where ρ_1 and ρ_2 govern the persistence of the state. We impose the usual assumptions on these two parameters to ensure stationarity.

There are now seven parameters to be estimated. We estimate these parameters by targeting the same eight moments described in the main text. As a result, our estimator is still an overidentified SMM estimator. The results are reported in Table D10.

All the parameters are precisely estimated and the model fits the empirical moments well. We estimate $\rho_1 > 0$ and $\rho_2 < 0$, indicating that AR(2) dynamics can account for some of the offsetting revisions and observed overreactions in the data. With that said, we note that controlling for adjacent revisions, there is still evidence of offsetting revisions over longer horizons. While such patterns cannot arise with an AR(2) process, they can arise under

[27] In this unreported exercise, we considered AR(2), AR(4), ARMA(1,1), ARMA(2,1) and ARMA(2,2) models.

Table D10: Model Estimation Results under AR(2) Dynamics

Panel A: Parameter Estimates			
	Parameter	Estimate	Standard error
Persistence of latent state	ρ_1	0.699	0.057
Persistence of latent state	ρ_2	-0.203	0.077
State innovation dispersion	σ_w	2.193	0.259
Public signal noise	σ_e	1.813	0.157
Private signal noise	σ_v	0.454	0.285
Probability of quarterly update	q	0.883	0.057
Probability of annual update	p	0.760	0.109
Panel B: Moments			
	Model	Data	t-statistics
Standard deviation of nowcast	2.390	2.399	-0.067
Correlation of nowcast with annual forecast	0.745	0.745	-0.061
Standard deviation of annual forecast	1.617	1.622	-0.056
Standard deviation of revision	2.012	1.999	0.103
Correlation of revision with lagged error	0.145	0.147	-0.035
Standard deviation of lag error	2.122	2.128	-0.068
RMSE nowcast	2.147	2.150	-0.030
RMSE annual forecast	1.477	1.477	0.000

Note: Panel A reports the model vs. data moments with t-statistics reported in the fourth column. Panel B reports the model parameters with point estimates reported in the third column and standard errors reported in the fourth column.

annual smoothing and temporal consistency.

Moreover, estimates of the state innovation and public noise variances are similar to those in Table 4. The private signal noise variance and the quarterly updating probability are estimated to be lower than the baseline estimates, while the annual updating probability is estimated to be higher. Relative to Table 7, these estimates imply a roughly similar level of information rigidity in quarterly real GDP forecasts, 0.152, and a meaningfully lower level of information rigidity in annual real GDP forecasts, 0.269. While the scope for overreactions declines in the AR(2) model, it is still capable of reproducing about 60% of the overreactions generated in the AR(1) model.