

Asymmetric Attention[†]

By ALEXANDRE N. KOHLHAS AND ANSGAR WALTHER*

We document that the expectations of households, firms, and professional forecasters in standard surveys simultaneously extrapolate from recent events and underreact to new information. Existing models of expectation formation, whether behavioral or rational, cannot account for these observations. We develop a rational theory of extrapolation based on limited attention, which is consistent with this evidence. In particular, we show that limited, asymmetric attention to procyclical variables can explain the coexistence of extrapolation and underreactions. We illustrate these mechanisms in a microfounded macroeconomic model, which generates expectations consistent with the survey data, and show that asymmetric attention increases business cycle fluctuations. (JEL C53, D83, D84, E23, E27, E32)

Given the central role of people's expectations in economics, it is important to have a theory of expectations formation that is consistent with the data. There is reason to believe that such a theory needs to be richer than the benchmark model of *full information* and *rational expectations*. Indeed, the original proponents of rational expectations were aware of this prospect. Muth (1961) allowed for “under-discounting” in his theory, noting that people may extrapolate from current events. Lucas (1972) studied agents who observe imperfect, noisy information, and later argued that “for most agents [...] there is no reason to specialize their information systems for diagnosing general movements correctly” (Lucas 1977, p. 21).

Many recent advances in the theory of expectations formation fall into one of two frameworks. On one hand, the noisy rational expectations approach proposed by Lucas has returned to popularity following the work of Woodford (2001) and Sims (2003). On the other hand, a common view is that such rational models cannot account for people's pervasive tendency to extrapolate from recent events, which

*Kohlhas: Institute for International Economic Studies (email: alexandre.kohlhas@iies.su.se); Walther: Imperial College London (email: a.walther@imperial.ac.uk). Mikhail Golosov was the coeditor for this article. First draft: January, 2018. We are indebted to three anonymous referees, George-Marios Angeletos, Ryan Chahrour, Martin Ellison, Nicola Gennaioli, Per Krusell, Kristoffer Nimark, David Stromberg, Robert Ulbricht, Laura Veldkamp, Mirko Wiederholt, and conference and seminar participants at the AEA Annual Meeting 2019, Barcelona Summer Forum, Cambridge University, Cornell University, CEPR Summer Symposium (Gerzensee) 2017, HKUST, Mannheim University, NORMAC 2017, Oxford University, Salento Summer Meetings, and SED 2018 for their comments. This research received financial support from the Lamfalussy Fellowship and Ragnar Soderberg Stiftelsen.

[†]Go to <https://doi.org/10.1257/aer.20191432> to visit the article page for additional materials and author disclosure statements.

has been documented in the survey data.¹ The latter view favors behavioral models of expectation formation that are consistent with extrapolation. The tension between these two frameworks is important, because the outcomes and dynamics of models with behavioral biases may differ from those with noisy rational expectations. Despite the obvious importance of this issue, no consensus has been reached.

In this paper, we argue that many existing models of expectation formation, whether behavioral or rational, cannot account for the survey evidence. This is because they cannot account for the fact that *overreactions* to recent events (i.e., extrapolation) often coincide with the type of *underreactions* to average new information that have been pointed out by Coibion and Gorodnichenko (2015). Our main contribution is to propose a unified model of expectation formation based on noisy rational expectations that resolves the friction between theory and data, and to explore its business cycle implications.

To empirically motivate our work, we demonstrate simultaneous overreactions and underreactions in a range of survey data.² The participants of standard surveys, reporting their expectations about future output and inflation, not only extrapolate from recent conditions, but also underreact to average information (as measured by average forecast revisions).

We show that a popular class of models, in which agents process signals of a forecasted variable (output, for concreteness), are inconsistent with such *simultaneous* over- and underreactions. This class includes standard behavioral models of extrapolation bias (e.g., Cutler, Poterba, and Summers 1990; Barberis et al. 2016), simple models of noisy rational expectations as derived from models of rational inattention (e.g., Sims 2003), as well as models that combine extrapolation bias or overconfidence with the presence of noisy information (e.g., Daniel, Hirshleifer, and Subrahmanyam 1998; Bordalo, Gennaioli, and Shleifer 2018). Intuitively, noisy information (or inattention) generates underreactions to new information, because individuals shrink their forecasts towards prior beliefs when the signals they observe are noisy. By contrast, extrapolation bias or overconfidence generates overreactions. We show that, on balance, when agents process signals of the forecasted variable, only one of these forces can dominate. In addition, we find that the same result extends to several influential models with a richer information structure (e.g., Lucas 1973; Lorenzoni 2009; Maćkowiak and Wiederholt 2009; Angeletos, Collard, and Dellas 2018). This is inconsistent with the simultaneous over- and underreactions that we find in the survey data.

Our core contribution is to develop a theory of extrapolation that is based on rational updating. We consider a model of forecasters who observe noisy information due to their limited attention. The distinguishing feature of our model is that forecasters do not passively observe noisy signals of aggregate conditions. Instead, agents observe noisy information of the various, structural components that comprise output, and can choose how much attention to pay to each component. The

¹ See, for example, Barberis et al. (2016); Bordalo, Gennaioli, and Shleifer (2018); and the references therein.

² Specifically, in Section I, we consider output and inflation forecasts from four of the most commonly used surveys on expectations: the Federal Reserve Bank of Philadelphia Survey of Professional Forecasters (SPF), the European Central Bank's (ECB) SPF, the Michigan Survey of Consumers, and the Livingston Survey.

combination of rational updating and noisy information implies that our theory remains consistent with observed underreactions.

In our model, output is the sum of several components. For example, these components could represent different inputs into the economy's production function, different sectors of the economy, or different variables in the economy's dynamic Euler equation for output. A population of forecasters observes a vector of noisy signals, where each signal contains information about a particular component. We think of *attention* to each component as the precision of the associated signal. Importantly, attention can be higher for some components than for others. We say that attention is *asymmetric* if agents receive a relatively more precise signal about some components. In this environment, we derive two main results.

The first main result is that asymmetric attention can explain the coexistence of extrapolation and underreactions, as long as attention centers on *procyclical* components. Consider an economy in which output is driven by only two components, which differ in their behavior over the business cycle. The first component is procyclical, while the second is countercyclical. Suppose that agents pay more attention to the procyclical component. Then, compared to the full-information benchmark, agents become more optimistic in booms and more pessimistic in busts, even though they adhere to Bayes' rule. As a result, the measured *overreactions* to recent output in the survey data can be viewed as an outcome of *underreactions* to countercyclical components. In addition, as long as agents' attention to the procyclical component remains imperfect, they still exhibit underreaction to new information on average, due to their rationally muted responses to noisy information. We extend this reasoning to a canonical forecasting problem with an arbitrary number of components. An auxiliary proposition generalizes our results to a comprehensive class of linear models.

Our second main result concerns the possible sources of asymmetric attention. In principle, asymmetric attention could arise from behavioral heuristics or salience effects (Gabaix 2017). Notwithstanding such alternatives, we show that asymmetric attention arises naturally in a rational framework, in which agents optimally choose how to allocate costly attention. With standard attention cost functions, agents in our framework find it optimal to pay asymmetric attention to components that are either particularly volatile or important for their decision-making. For example, consider a firm that reports its expectation about future output. In line with the conclusions in Lucas (1977), this firm has an incentive to focus its attention on the components of output that correlate closely with its own local conditions, especially if these components are also volatile. Coibion, Gorodnichenko, and Kumar (2018b) uses detailed firm-level data to provide direct evidence of firms' incentive to pay asymmetric attention to volatile and important variables.

Combining our two results, we conclude that a rational model of limited attention can simultaneously explain extrapolation and underreaction to aggregate information, as long as the volatile or important components of output that attract attention are also procyclical. This connects our results to those of Woodford (2001), Nimark (2008), and Angeletos and Huo (2021), among others, which argue that limited attention can account for the myopia and anchoring to past outcomes often documented in macroeconomics. We demonstrate that models of limited attention also have the potential to be consistent with extrapolation.

We show that an additional testable implication of our explanation, in terms of the aggregate data, is that expectations should be more precise than pure time-series forecasts (e.g., forecasts from autoregressive integrated moving average (ARIMA) models). Consistent with this prediction, we update estimates from Stark (2010) to show that forecasters' survey expectations of output growth consistently outperform simple time-series models, especially at short horizons.

To explore the implications of our framework, and to provide an example of the sources of asymmetric attention, we apply our framework to a standard macroeconomic model with flexible prices in the spirit of Angeletos, Iovino, and La'O (2016). In the model, firms choose output under imperfect information about productivity. We show that, in equilibrium, firms' output choices can be split into two components: (i) firm beliefs about a *productivity component*, which reflects their own productivity; and (ii) firm beliefs about an *aggregate supply component*, which summarizes the equilibrium effect of other firms' choices on individual firm output. Maćkowiak and Wiederholt (2009) proposes a closely related decomposition. When we sum across firms, aggregate output thus becomes the simple sum of the two components.

We show that, for standard parameter values, two key conditions are satisfied: First, the productivity component is procyclical, while the aggregate supply component is countercyclical. The latter follows because economy-wide expansions tend to increase firms' costs, leading each individual firm to reduce its output relative to its partial equilibrium choice. Second, if attention is costly, firms optimally choose to pay asymmetric attention to their own productivity, because this component is substantially more volatile. As a result of these two conditions, and in line with our two main results, firms' expectations of future aggregate output exhibit both extrapolation and underreactions to recent forecast revisions, relative to the full information benchmark. This is qualitatively consistent with the survey evidence. The model also fits the empirical size of these effects well.

We use the macroeconomic model to explore the business cycle implications of firms' asymmetric attention choices. We show that asymmetric attention to local components leads to more persistence and volatility in aggregate output than an equivalent model with symmetric attention. We further document that the calibrated model exhibits an increase in extrapolation post-Great Moderation, and argue that firms' optimal attention choices may have contributed to the increased persistence of output during this period.

Finally, two wider implications of our analysis are worth noting. First, in the tradition of Lucas (1977), our macroeconomic model focuses on a lack of attention to equilibrium effects as the driver of extrapolation. As such, our results speak to a literature in behavioral finance, which models the neglect of equilibrium effects as fundamental behavior, and uses this to account for investment patterns (e.g., Greenwood and Hanson 2015).

Second, motivated by the survey evidence, we focus on a setting in which agents' forecasts appear to overreact to a particular type of public information (i.e., recent realizations of the forecasted variable). However, as we illustrate, a model of asymmetric attention may be equally consistent with underreactions to other types of public information, depending on how this information correlates with the variables

to which agents pay attention.³ We therefore view this paper, more generally, as taking a first step towards integrating observed over- and underreactions to new information into a unified, rational framework.

Related Literature.—In addition to the literature cited above, this paper relates to four areas of research. We review these in reverse chronological order, starting with the most recent and ending with the long history of thought on extrapolative and adaptive expectations.

First, our paper reconciles overreactions to recent outcomes of the forecasted variable with underreactions in *average forecast revisions*. In contemporaneous and closely related work, Bordalo et al. (2020) proposes a behavioral model that can reconcile similar underreactions to *average* forecast revisions with overreactions to *individual forecast revisions* (see also Fuhrer 2017 and Broer and Kohlhas 2019). Their notion of overreactions is distinct from the overreactions to current *aggregate conditions*, such as current output, on which we focus in this paper.⁴ As we demonstrate in Section I, simple versions of the framework in Bordalo et al. (2020) cannot account for the simultaneous occurrence of underreactions to average revisions and overreactions to current output growth that we document in the data. Furthermore, we show that an extended version of our model can fit the stylized facts in both papers. Crucially, this would not be possible in a model with symmetric attention choices. In another recent paper, Angeletos, Huo, and Sastry (2020) shows how a combination of different behavioral biases can generate the empirical estimates in our work, as well as the estimates in Coibion and Gorodnichenko (2015) and Bordalo et al. (2020). We view this recent strand of research as presenting related and complementary steps towards a unified model of expectations that is consistent with over- and underreactions to new information.

Second, in common with a vast literature in macroeconomics since Lucas (1972), we emphasize the importance of imperfect information for business cycle dynamics. Prominent studies, among many others, are Woodford (2001); Mankiw and Reis (2002); Lorenzoni (2009); Blanchard, L’Huillier, and Lorenzoni (2013b); Angeletos and La’O (2013b); Maćkowiak and Wiederholt (2015); and Chahrour and Ulbricht (2018). We emphasize the role of agents who optimally choose how to allocate their scarce attention, and we build on the complementary literatures on “optimal information choice” (e.g., Veldkamp 2011; Hellwig, Kohls, and Veldkamp 2012) and “rational inattention” (e.g., Sims 2003, Maćkowiak and Wiederholt 2009, Wiederholt 2010). The contribution of our paper, in this context, is to highlight that models of imperfect information can also be consistent with the observed overreactions in the survey data.

Third, we leverage the existing evidence on survey expectations. Pesaran (1987) summarizes the early evidence on deviations from full information and rational expectations, and Zarnowitz (1985) shows that survey data is consistent with models of noisy, private (instead of common, perfect) information. Relatedly, Ehrbeck and Waldmann (1996) explores the sources of bias in professional forecasts and

³Underreactions to public information are documented, for example, in Barberis, Shleifer, and Vishny (1998); Daniel, Hirshleifer, and Subrahmanyam (1998). Eyster, Rabin, and Vayanos (2019) reviews further related evidence.

⁴We discuss the relationship between our work and that of Bordalo et al. (2020) in detail in Section IID.

conclude that these are unlikely to derive from agency-based considerations. More recently, Coibion and Gorodnichenko (2012b, 2015) demonstrate underreactions to average forecast revisions (see also Andrade and Le Bihan 2013, and Fuhrer 2017), which form part of the motivation for this paper.

Finally, our focus on overreactions to recent outcomes connects this paper to the literature on adaptive and extrapolative beliefs. This includes the early work of Goodwin (1947), Cagan (1956), and Muth (1961), the experimental work on the psychology of subjective probabilities as explored by Kahneman and Tversky (1972) and Andreassen and Kraus (1988), and the modern treatments of extrapolation by de Long et al. (1990); Cutler, Poterba, and Summers (1990); Fuster, Hebert, and Laibson (2012); Greenwood and Shleifer (2014); Barberis et al. (2016); and Bordalo, Gennaioli, and Shleifer (2018). This paper is the first, to our knowledge, to combine the empirical insights of this literature with a model that can also generate underreactions to aggregate expectations.

I. Motivating Evidence and Existing Theory

In this section, we revisit two simple tests of full information and rational expectations. We document a new stylized fact: participants' expectations in standard surveys *simultaneously* overreact to recent realizations of the forecasted variable (i.e., extrapolate from recent events), but underreact in their forecast revisions. We then derive the predictions of a popular set of existing models and argue that these models cannot account for this observation.

A. Simultaneous Over- and Underreactions

We start by considering forecasts of US output growth from the *Survey of Professional Forecasters* (SPF).⁵ The SPF is a survey of between 20 and 100 professional forecasters and is conducted quarterly by the Federal Reserve Bank of Philadelphia. Real GDP/GNP growth estimates are available from 1968:IV at a quarterly frequency. We focus on output forecasts for two reasons. First, because expectations about future output play a central role in the economy as determinants of consumption, inflation, and asset prices. Second, because data on output forecasts are available for a longer time-span than forecasts of most other variables. We later explore the robustness of our empirical estimates by considering forecasts of inflation, as well as alternative survey datasets for the United States and the euro area.

We let y_{t+k} denote year-on-year output growth at time $t + k$. Consider a survey with respondents indexed by $i \in \{1, 2, \dots, I\}$, and let $f_{it}y_{t+k}$ denote the forecast of y_{t+k} reported by survey respondent i at time t . The respondent's *forecast error* is $y_{t+k} - f_{it}y_{t+k}$. A negative forecast error thus corresponds to an overestimate of y_{t+k} . A well-known implication of *full information and rational expectations* (FIRE) is that individual forecast errors should be unpredictable. Under FIRE, no variable that

⁵ The SPF is the oldest quarterly survey of individual macroeconomic forecasts in the United States, dating back to 1968. The SPF was initiated under the leadership of Arnold Zarnowitz at the ASA and the NBER, which is why it is also still often referred to as the ASA-NBER Quarterly Economic Outlook Survey (Croushore 1993).

is observable at time t should correlate with $y_{t+k} - f_{it}y_{t+k}$. We rely on two common tests of this prediction.

The first test is a regression of forecast errors on current output growth,

$$(1) \quad y_{t+k} - f_{it}y_{t+k} = \alpha_i + \gamma y_t + \xi_{it},$$

where α_i is a constant, which also captures individual fixed effects, and ξ_{it} is an error term. The second test is a regression of forecast errors on average forecast revisions,

$$(2) \quad y_{t+k} - f_{it}y_{t+k} = \alpha_i + \delta(\bar{f}_t y_{t+k} - \bar{f}_{t-1} y_{t+k}) + \xi_{it}.$$

The term $\bar{f}_t y_{t+k} - \bar{f}_{t-1} y_{t+k}$ on the right-hand side is the average change in respondents' forecasts when they are asked twice (at dates $t-1$ and t) to forecast the same future realization y_{t+k} . A positive revision arises when good news about future output arrives between $t-1$ and t . This specification closely follows the test proposed by Coibion and Gorodnichenko (2015).⁶

The prediction of the FIRE benchmark is that the coefficients γ and δ in (1) and (2) should both be zero, because both current output growth and the latest forecast revision are observable at time t . It is useful to note that both (1) and (2) are tests of the *joint hypothesis* of full information and rational expectations. A rejection of the FIRE prediction reveals *either* that forecasters are reporting irrational expectations, *or* that they have imperfect information about current output (for $\gamma \neq 0$) or average forecast revisions (for $\delta \neq 0$). However, it does *prima facie* not reveal which one of these hypotheses is rejected.

The raw data already hint at deviations from the FIRE benchmark. Figures 1 and 2 plot average one-year-ahead forecast errors (the average left-hand side of (1) and (2) across respondents, with $k = 4$) over time and compare them, respectively, to current realizations of output growth (the right-hand side of (1)) and average one-quarter revisions (the right-hand side of (2)).⁷ In Figure 1, forecasts are frequently overoptimistic, with associated negative forecast errors when current output growth is high, and vice versa when current growth is low. This suggests that respondents extrapolate from recent events; agents are systematically too optimistic in booms and too pessimistic in busts. Figure 2, by contrast, suggests that forecast errors and average forecast revisions are positively correlated within our sample. All else equal, this indicates that agents underreact to new information on average, as they are too pessimistic after positive forecast revisions, and vice versa after negative revisions.

⁶Coibion and Gorodnichenko (2015) use *average* forecast errors $y_{t+k} - \bar{f}_{it}y_{t+k}$ as the dependent variable in (2). We prefer the individual-level regression because it is easier to compare its results to candidate theories of individual expectation formation, and also because it allows for respondent-level fixed effects and assigns equal weight to all individual forecasts in an unbalanced panel such as ours. For completeness, we report both average- and individual-level estimates throughout the paper and the online Appendix.

⁷We use real-time data to measure current realizations of output growth. Because the response deadline for the SPF is about one-week from the US Bureau of Economic Analysis's (BEA) first release of output growth, we on the right-hand side of (2) average this release's value with its previous quarter's realization. This is to precisely capture the current conditions at the time the respondent institutions determine their published forecast (e.g., Croushore and Stark 2019; Bordo, Gennaioli, and Shleifer 2018). We do not make this adjustment for other variables and datasets that we consider below, as for these there is time to include information into published forecasts. Table C.6 in the online Appendix shows that our results are similar using either of the two quarters' output growth values.




FIGURE 1. OVERREACTIONS IN OUTPUT GROWTH FORECASTS

Notes: Mean one-year-ahead forecast error of output growth from the SPF on the left vertical axis, and the current realization on the right axis. Both scales are in percent year on year. Current realizations are measured as the average of the BEA's first release value and its previous quarter's realization. This is to account for the timing of the SPF survey (see footnote 7 for further discussion).




FIGURE 2. UNDERREACTIONS IN OUTPUT GROWTH FORECASTS

Notes: Mean one-year-ahead forecast error of output growth from the SPF on the left vertical axis, and the one-quarter revisions on the right axis. Both scales are in percent year on year.

Table 1 confirms these impressions and reports estimates of (1) and (2) using the SPF data on one-year-ahead forecasts ($k = 4$). In the first column, we estimate (1) and find that γ is negative and statistically significant. This once more suggests

TABLE 1—ESTIMATED OVER- AND UNDERREACTIONS IN THE SPF

	(1)	(2)	(3)
<i>Panel A. Individual forecast error</i>			
Current realization	−0.12 (0.05)	—	−0.14 (0.04)
Average revision	—	0.68 (0.19)	0.71 (0.18)
Observations	7,104	7,065	7,008
<i>F</i> -statistic	169.2	449.6	363.8
<i>R</i> ²	0.02	0.06	0.10
<i>Panel B. Average forecast error</i>			
Constant	0.02 (0.19)	−0.09 (0.10)	0.25 (0.15)
Current realization	−0.10 (0.05)	—	−0.13 (0.05)
Average revision	—	0.78 (0.26)	0.84 (0.25)
Observations	196	195	194
<i>F</i> -statistic	3.29	16.6	11.9
<i>R</i> ²	0.02	0.08	0.11

Notes: Panel A: estimates of regressions (1) and (2) with individual (respondent) fixed effects. The top and bottom 1 percent of forecast errors and revisions have been removed. Table C.1 in the online Appendix shows similar results without removing outliers. Double-clustered robust standard errors in parentheses. Panel B: estimates of regressions (1) and (2) with average forecast errors $y_{t+k} - \bar{f}_t y_{t+k}$ as the left-hand side variable. Robust standard errors in parentheses. Current realizations are measured as the average of the BEA's first release value and its previous quarter's realization. Sample: 1970:IV–2019:IV.

extrapolation, or *overreactions* to recent realizations of output growth. In the second column, we estimate (2) using one-quarter average revisions. We find that δ is positive and significant, which is consistent with average forecast revisions *underreacting* to overall new information received within the period.

The third column confirms these results in a multiple regression. The multivariate estimates are similar to those in the univariate case. This suggests that the univariate results are not biased by correlation between output realizations and forecast revisions.

Taken individually, the over- and underreactions documented in Table 1 are in line with previous estimates. Bordalo, Gennaioli, and Shleifer (2018), for example, report evidence on extrapolation based on the average-level version of regression (1). For regression (2), our estimates update those reported by Coibion and Gorodnichenko (2015, Figure I). Our results demonstrate that, in addition, extrapolation and underreactions occur *simultaneously* in the SPF data.

In contemporaneous and closely related work, Bordalo et al. (2020) analyzes a different type of “overreactions” in survey expectations to that documented in Table 1. Specifically, Bordalo et al. (2020) analyzes overreactions to individual forecast revisions. By contrast, we use regression (1) to emphasize overreactions to recent realizations of the forecasted variable. For now, we continue to focus on our

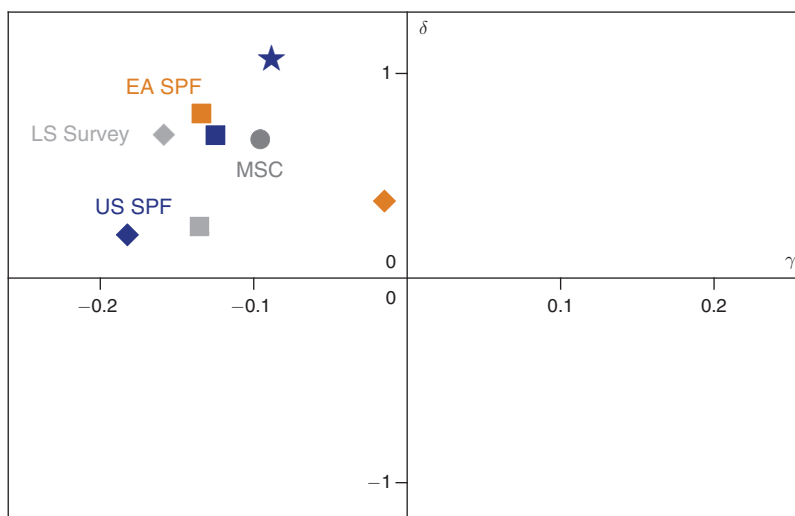


FIGURE 3. ESTIMATED OVER- AND UNDERREACTIONS ACROSS SURVEYS

Notes: Estimates of the coefficients γ and δ from (1) and (2) using individual forecast errors $y_{t+k} - \hat{f}_{it,t+k}$ as the dependent variable. *US SPF* represents the estimates for the US Survey of Professional Forecasters, *EA SPF* the ECB's Survey of Professional Forecasters, *LS Survey* the Livingston Survey, and lastly *MSC* the Michigan Survey of Consumers. \square = GDP forecasts, \diamond = consumer price index (CPI) inflation forecasts, \star = GDP deflator inflation forecasts, and \circ = MSC CPI inflation forecasts that have been instrumented. All estimates are for one-year-ahead forecasts with the exception of the Livingston Survey (three-quarters ahead). Estimates of (2) use semiannual revisions (Livingston Survey), annual revisions (EA SPF), or one-quarter revisions (others). Figures C.1 and C.2 in the online Appendix illustrate the robustness of the estimates above to alternative sample assumptions and the use of average forecast errors as the dependent variable.

regression (1). In Section IID, we provide a detailed discussion of these distinct notions of overreaction.

We obtain similar estimates to those in Table 1 beyond forecasts of output growth in the US SPF. Figure 3 summarizes estimates of (1) and (2) for output and inflation forecasts from the *euro area SPF*, the *Livingston Survey* (which covers academic institutions, investment banks, nonfinancial firms, and government agencies), and the *Michigan Survey of Consumers*.⁸

We plot the coefficient γ on current realizations in (1) on the horizontal axis in Figure 3, and the coefficient δ on average forecast revisions in (2) on the vertical axis.⁹

⁸The Livingston Survey is a semiannual survey, which collects one, three, and five quarter-ahead forecasts of several macroeconomic variables (Croushore 1997). The Michigan Survey of Consumers contains consumers' inflation forecasts (Curtin 1982). A drawback of the monthly Michigan Survey of Consumers is that only one-year ahead forecasts of consumer price inflation are available. Revisions to forecasts at a fixed horizon cannot be constructed. To estimate (2), we therefore follow Coibion and Gorodnichenko (2015) and replace *ex ante forecast revisions* with the quarterly *ex ante forecast changes* and instrument this variable with the (log) oil price change. This approach provides an asymptotically consistent estimate. The euro area's SPF (Garcia 2003) collects the same information as the US SPF.

⁹Some of our estimates of (2) are direct updates of estimates reported by Coibion and Gorodnichenko (2015) using average forecast errors as the dependent variable. In particular, Coibion and Gorodnichenko (2015) also reports estimates of (2) using CPI inflation forecasts from the Livingston Survey and the Michigan Survey of Consumers, GDP deflator inflation forecasts from the US SPF, as well as inflation forecasts from the euro area (although from the Consensus Economic Survey and not the euro area SPF). All of these estimates are comparable to ours. Relative to their work, we focus on *simultaneous* estimates of (2) and (1), and cover a wider range of data sources for output growth forecasts, which are the focus of our analysis.

All of our estimates fall into the upper-left quadrant of the figure, where we simultaneously find that $\gamma < 0$ (overreaction) and that $\delta > 0$ (underreaction). Table C.7 in the online Appendix contains the associated regression results. Specifically, with the exception of the euro area and Michigan CPI inflation forecasts, and the GDP deflator forecasts from the US SPF, all overreaction coefficients in Figure 3 are statistically significant at the 5 percent level.

Tables C.2–9 in the online Appendix contain further robustness checks. We show that simultaneous over- and underreactions extend to multivariate versions of (1) and (2), to the use of average forecast errors $y_{t+k} - \bar{f}_t y_{t+k}$ as the dependent variable, and to different forecast horizons,¹⁰ timing conventions, and assumptions about trends in the data. We also split the sample and find similar patterns in the post-1992 sample (to account for any potential structural break in the inflation series)¹¹, as well as both pre- and post-Great Moderation.

Finally, we also consider two alternative tests from the literature to confirm the robustness of our results. First, following Coibion and Gorodnichenko (2015), we report estimates of the unconstrained version of (2) with potentially different coefficients on $\bar{f}_t y_{t+k}$ and $\bar{f}_{t-1} y_{t+k}$ (Table C.5). We fail to reject the null hypothesis that the coefficients sum to zero, validating the specification in (2). Second, in online Appendix D, we consider the projection of average forecast errors and current output growth on identified productivity shocks, as in Coibion and Gorodnichenko (2012b). Consistent with underreactions, we find a positive correlation between the conditional response of forecast errors and the response of output growth.

In summary, the results in Table 1 and Figure 3 document systematic overreactions to recent realizations of the forecasted variable (i.e., extrapolation), but *simultaneous* underreactions to average forecast revisions. This clearly constitutes a rejection of the joint hypothesis of full information and rational expectations. In the next subsection, we consider a range of existing models that relax either full information or rational expectations. We argue that one can also use our stylized facts to determine whether existing alternative theories of expectation formation are consistent with the data.

B. Existing Theories of Expectation Formation

We compare our estimates to a parsimonious framework, where agents observe noisy signals of the forecasted variable, which captures several popular models of expectation formation. On the one hand, we show that rational forecasts are inconsistent with overreactions to current output (i.e., $\gamma < 0$ in (1)), and that this extends to a collection of richer models. On the other hand, we show that several popular behavioral alternatives, which are able to generate $\gamma < 0$, cannot simultaneously generate underreactions to average information (i.e., $\delta > 0$ in (2)).

¹⁰The point estimates with shorter forecast horizons decline in magnitude and significance. This is consistent with a greater importance of noise in shorter horizon forecasts (Coibion and Gorodnichenko 2015). Table C.3 shows estimates after detrending output growth. We include this table for completeness, despite the potential concern that the detrending operation uses information from the whole sample, introducing look-ahead bias. However, regardless of this issue, the detrended data yield very similar estimates to the raw data.

¹¹The Federal Reserve Bank of Philadelphia took over ownership of the SPF in 1990:II.

Consider a continuum of measure one of agents who make forecasts of future output y_{t+k} . We assume that output y_t follows the autoregressive process:

$$(3) \quad y_t = \rho y_{t-1} + u_t, \quad u_t \sim \mathcal{N}(0, \sigma_u^2),$$

where $\rho \in (0, 1)$ and u_t is serially uncorrelated. We focus on the case in which output follows a stationary process, and on agents' forecasts about the level of output. We make this choice to simplify the exposition of what follows, and for consistency with the existing literature in macroeconomics. With minor modifications, our theoretical results also extend to the case where, as in our data, the level of output has potential unit roots or time trends, and agents make forecasts about its growth rate.

At the start of each period, each agent $i \in [0, 1]$ observes a noisy signal of current output,

$$(4) \quad z_{it} = y_t + \epsilon_{it}, \quad \epsilon_{it} \sim \mathcal{N}(0, \sigma_\epsilon^2),$$

where the noise in agents' signals ϵ_{it} is independent of u_t at all horizons with $\text{cov}[\epsilon_{it}, \epsilon_{js}] = 0$ for all $i \neq j$ and $t \neq s$. We write $\Omega_{it} = \{z_{is}\}_{s \leq t}$ for agent i 's information set at date t .¹²

We assume that agents' forecasts follow a *recursive forecast equation*, which generalizes the textbook Kalman filter. Let $f_{it}y_{t+k}$ and $f_{it-1}y_{t+k}$ denote agent i 's forecasts of future output at dates t and $t-1$, respectively, and let $f_{it-1}z_{it}$ be her forecast of her own signal one period ahead. Agent i 's output forecast then follows the updating equation:

$$(5) \quad f_{it}y_{t+k} = \lambda f_{it-1}y_{t+k} + g_k(z_{it} - \lambda f_{it-1}z_{it}),$$

where $f_{it}y_{t+k} = \rho^k f_{it}y_t$. As with the textbook Kalman filter, the agent starts with her forecast of output at time $t-1$, and updates it in proportion to the new information in her signal at time t . Departing from the standard filter, we allow $g_k \geq 0$ to be an arbitrary gain parameter that measures agents' responsiveness to new information.¹³ We also allow the prior update parameter $\lambda \in [0, 1]$ to be less than one. Despite its simplicity, the formulation in (5) nests a wide range of existing models of expectation formation. We demonstrate this through a series of examples, which we delineate into rational and behavioral theories.

Noisy Rational Expectations.—Agents' forecasts equal their conditional expectation $f_{it}y_{t+k} = E[y_{t+k} | \Omega_{it}]$ and follow (5) with a gain parameter $g_k = \text{cov}_{t-1}[y_{t+k}, z_{it}] / \text{var}_{t-1}[z_{it}] < \rho^k$ while $\lambda = 1$. This specification is identical to those from models with noisy rational expectations (Woodford 2001) or rational

¹²We allow agents to observe an infinite history of signals, so that their signal extraction problem is initialized in steady state at date 0. This assumption follows the convention in, e.g., Maćkowiak, Matějka, and Wiederholt (2018).

¹³To ensure that forecasts in (5) are well-defined, we impose that $g_k = \rho^k g_0$ with $g_0 \in (0, 2)$.

inattention (Sims 2003).¹⁴ The special case in which agents observe output without noise ($\sigma_\epsilon = 0$) corresponds to the case of FIRE, and implies that $f_{it}y_{t+k} = \rho^k y_t$ with $g_k = \rho^k > 0$.

Behavioral Expectations.—A common way to model behavioral biases is to assume agents perceive the data-generating process to be different from its true parametrization, but then update correctly under this wrong model. Equation (5) captures several of such cases:

- **Overconfidence.**—Agents overestimate the precision of new information. They believe that the variance of the noise in their signals is $\hat{\sigma}_\epsilon^2 < \sigma_\epsilon^2$ (e.g., Daniel, Hirshleifer, and Subrahmanyam 1998; Hirshleifer, Lim, and Teoh 2011).¹⁵ Agents forecasts follow the recursive forecast equation in (5) with a sensitivity parameter $g_k \in (0, 1)$ that exceeds its rational value and $\lambda = 1$.
- **Extrapolation.**—Agents overestimate the extent to which current output predicts future realizations. They observe output without noise ($\sigma_\epsilon^2 = 0$), but believe that the persistence parameter for output is $\hat{\rho} > \rho$ (e.g., de Long et al. 1990; Fuster, Hebert, and Laibson 2012). Agents' forecasts satisfy (5) with a sensitivity parameter $g_k = \hat{\rho}^k > \rho^k$ and $\lambda = 0$.¹⁶
- **Diagnostic Expectations.**—The model in Bordalo, Gennaioli, and Shleifer (2018) and Bordalo et al. (2020) corresponds to the overconfidence case, but the effect of overconfidence is temporary and does not affect forecasts at future dates. Equation (5) is replaced by $f_{it}y_{t+k} = E_{it-1}y_{t+k} + g_k(z_{it} - E_{it-1}y_t)$, where g_k exceeds its rational value. Despite the nonrecursivity of forecasts, we include the model in this list because the properties of its forecast errors $y_{t+k} - f_{it}y_{t+k}$ depend only on $(\rho^k - g_k)(y_{t+k} - E[y_{t+k}|\Omega_{it}])$, and thus exclusively on those from the noisy rational expectations case and (5) (Corollary 1 in Appendix AA).

We now characterize the results that an econometrician would obtain when estimating (1) and (2), assuming that the true data-generating process satisfies (3) to (5).

PROPOSITION 1: *Suppose agents form their expectations according to (5), based on signals in (4). Then, the coefficients γ in (1) and δ in (2) both have the same sign as $\rho^k - g_k$.*

Proposition 1 demonstrates that models described by the recursive forecast equation, such as the rational and two behavioral models above, all imply either underreactions in both of our main regressions ($\gamma > 0$ and $\delta > 0$), or overreactions ($\gamma < 0$ and $\delta < 0$). Indeed, Proposition 1 also implies that the coefficients γ and δ in the model of diagnostic expectations also both have the same sign as $\rho^k - g_k$ (Corollary 1 in Appendix AA). This is at odds with our empirical estimates of

¹⁴ A more comprehensive list of papers in this tradition is in the introduction. The Gaussian signal z_{it} we have specified is optimal in a rational inattention setting if agents minimize their squared forecast errors and their cost of processing information is based on the reduction in entropy (see Maćkowiak, Matějka, and Wiederholt 2018).

¹⁵ For further analysis of overconfidence, see Broer and Kohlhas (2019) and the references therein.

¹⁶ Thus, $f_{it}y_{t+k} = \hat{\rho}^k y_t$. The introduction contains a further list of references using such forecasts.

simultaneous over- and underreactions. One can see this discrepancy clearly in terms of Figure 3. Proposition 1 shows that an econometrician's estimates will fall either into the upper-right quadrant or the lower-left quadrant of the figure. This is inconsistent with our empirical estimates that center on the upper-left quadrant.

To interpret Proposition 1 further, recall that agents' gain parameter in the FIRE case, in which they perfectly observe current output, is equal to $g_k = \rho^k$ (since $f_{it}y_{t+k} = \rho^k y_t$). Proposition 1 states that there are two possible parametric regions, corresponding to systematic underreactions or overreactions, depending on whether agents' responsiveness to new information g_k is smaller or greater than in the FIRE benchmark.

Two counteracting effects determine the size of g_k . First, the presence of noise in agents' signals dampens agents' responsiveness to new information, which rationally pushes g_k below its FIRE value. This effect, all else equal, creates measured underreactions: An econometrician estimating (1) and (2) has access to more information than agents in the model, because he observes current output and average forecast revisions perfectly. As a result, forecast errors are predictable. And because agents respond to noisy information in a muted fashion, this predictability takes the shape of measured underreactions. Second, behavioral biases, such as overconfidence or extrapolation, heighten agents' responsiveness to new information, which in turn increases the gain coefficient. However, as Proposition 1 shows, only one of these forces can come to dominate the sufficient statistic g_k . Hence, in all of the cases above, agents either over- or underreact, but do not over- and underreact *simultaneously*.¹⁷

We conclude that a popular class of models, in which agents form Bayesian or non-Bayesian expectations based on noisy signals of the forecasted variable, is inconsistent with simultaneous over- and underreactions. In particular, it is clear that to explain the survey data, we must consider a model with more than one sufficient statistic for belief formation. In the next section, we achieve this aim by proposing a noisy rational expectation model in which agents pay limited but asymmetric attention to different structural components of the forecasted variable. Before turning to our model, we however briefly consider more sophisticated existing models of noisy rational expectations.¹⁸

We focus on richer models from two influential strands of literature. First, the literature on rational inattention includes more sophisticated models following Maćkowiak and Wiederholt (2009), in which agents rationally allocate their attention between aggregate and individual-specific conditions. Individual-specific conditions, and the signals that agents obtain about them, are uncorrelated with aggregate output by assumption. Hence, forecasts of future aggregate output behave *as if* agents obtained only a noisy signal of output

¹⁷For the same reason, a simple model with heterogeneous expectation formation among agents is also inconsistent with our estimates. In an economy with heterogeneous types of forecasters, who have different degrees of behavioral biases or limited attention, the generalized Kalman gain g_k in our formulation can be reinterpreted as the weighted average of each type's response to new information. Hence, average forecasts will either over- or underreact, but cannot do so at the same time.

¹⁸In addition, online Appendix E characterizes the more sophisticated behavioral model in Angeletos, Collard, and Dellas (2018), which introduces a small deviation from rational expectations into a model of dispersed information. Intuitively, agents in their model adjust their expectations in proportion to exogenous confidence shocks. We show that this model predicts *overreactions* to both output and average revisions (i.e., $\gamma < 0$, $\delta < 0$).

itself. Indeed, online Appendix E.1 shows that the above noisy rational expectations case, where $\gamma > 0$, exactly describes output expectations in Maćkowiak and Wiederholt (2009).

Second, we consider models with dispersed information in which agents observe local economic conditions (on “islands”) accurately but economy-wide conditions only with noise (e.g., Lucas 1973, Lorenzoni 2009). In online Appendix E.2 and E.3, we explicitly solve the models in Lucas (1973) and Lorenzoni (2009), and show that these models also generate underreactions to current output ($\gamma > 0$). The intuition is similar to that in the simple model with noisy observations of output: agents have less information about aggregates than the econometrician, and they respond to this information in a muted fashion, which creates underreactions. Indeed, we show that one can directly use (5) and the noisy rational expectations case above to obtain an analytical expression for the underreactions in Lucas (1973).

To summarize, it is instructive to view the results in this section in terms of our empirical findings using (1) and (2). Our estimates show that $\gamma < 0$ and $\delta > 0$, and reject the FIRE benchmark. This reveals that *either* the assumption of full information *or* the assumption of rationality is violated. However, our analysis of existing models establishes that it is not obvious how to match the data by relaxing either assumption. Although the list of models we have considered is not exhaustive, we are unaware of a pre-existing model that can explain our results. This motivates the development of our model in the next section.

II. Asymmetric Attention

In this section, we consider a rational model of limited attention. The central difference to the standard model from the previous section is that we view output as comprised of a set of structural components. We show that the over- and underreactions that we have documented can be rationalized if agents pay more attention to some components than others; that is if agents’ attention is *asymmetric*. Our approach in this section is to take attention choices as given and derive conditions under which the model can account for our empirical results. In the next section, we then examine the possible sources of asymmetric attention.

A. Environment

A continuum of measure one of agents are asked to forecast future output y_{t+k} . Aggregate output y_t is driven by the sum of N structural components x_{jt} ,

$$(6) \quad y_t = x_{1t} + x_{2t} + \cdots + x_{Nt}.$$

These components could, for example, represent different inputs into the economy’s production function, different sectors of the economy, or different variables in firms’ optimal production plans. We discuss one such example at length in Section IV. Each component x_{jt} is determined by the linear relationship

$$(7) \quad x_{jt} = a_j \theta_t + b_j u_{jt}, \quad u_{jt} \sim \mathcal{N}(0, 1),$$

where θ_t denotes a latent factor that follows the autoregressive process

$$(8) \quad \theta_t = \rho\theta_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \tau_\eta^{-1}),$$

with $\rho \in (0, 1)$. The error terms u_{jt} and η_t are serially uncorrelated, mutually independent, and it is common knowledge that $\theta_0 \sim \mathcal{N}(0, \tau_\theta^{-1})$. As a result, each component depends both on the common latent factor θ_t and on a transitory, component-specific shock u_{jt} .

The output response to a positive fundamental shock $d\theta_t > 0$ is $dy_t/d\theta_t = \sum_j a_j$. We assume that $\sum_j a_j > 0$ without loss of generality, so that output correlates positively with θ_t . The contribution of component x_{jt} to this output response is a_j . We refer to a component x_{jt} as *procyclical* if $a_j > 0$, so that x_{jt} reinforces the response of output to the latent factor. Analogously, we say that x_{jt} is *countercyclical* if it dampens the response with $a_j < 0$.

Output and its components are not directly observable to agents, because of their limited attention. Instead, each agent $i \in [0, 1]$ observes the history of N noisy signals

$$(9) \quad z_{ijt} = x_{jt} + q_j \epsilon_{ijt}, \quad \epsilon_{ijt} \sim \mathcal{N}(0, 1), \quad j = \{1, 2, \dots, N\},$$

where q_j parameterizes the noise (or inattention) in agents' signals about the j th component, and ϵ_{ijt} is an idiosyncratic error term. Agent i 's information set at time t is the history of her past signals $\Omega_{it} = \{z_i^0 \cup (z_{i1s}, \dots, z_{iNs})_{s=1}^t\}$.¹⁹ Agents thus infer information about the latent factor θ_t from noisy signals of x_{jt} that may covary either positively ($a_j > 0$) or negatively ($a_j < 0$) with the latent factor itself.

Notice that there are two key differences between this environment and that in Section I, which also nested a rational case with noisy signals. First, output is determined by several underlying components. Second, agents learn about these components separately: the information structure in (9) restricts agents to observing conditionally independent signals of each component. This formalizes the idea that paying attention to one component is a separate activity from paying attention to another. Combined, these features capture the notion that, to form expectations, individuals first need to pay attention to information about the various components of the forecasted variable, and then combine these different pieces of information into a single prediction. The conditional independence embedded in (9), combined with a component-based structure in (6), is a simple and common way to model this idea (see, e.g., Maćkowiak and Wiederholt 2009). We discuss the role of these restrictions in more detail in Section III, where we also consider an alternative setup with fully flexible information design.

¹⁹We assume that in the initial period $t = 0$, all agents receive an (infinitely) long sequences of signals from (9) of the N components, denoted by z_i^0 . This assumption follows the convention in the literature (see, e.g., Maćkowiak, Matějka, and Wiederholt 2018). By allowing agents to observe an infinite history of signals initially, we ensure that their signal extraction problem is initialized in steady state.

B. Definition of Attention

To characterize agents' attention to the various structural components, we transform the noise parameters q_j in (9) into the normalized parameters

$$(10) \quad m_j \equiv \frac{\text{var}(x_{jt}|\theta_t)}{\text{var}(z_{ijt}|\theta_t)} = \frac{b_j^2}{b_j^2 + q_j^2} \in (0, 1).$$

These parameters measure the sensitivity of agents' expectations to new information about the j th structural component. Suppose that agent i knows θ_t , and is then asked to predict component x_{jt} based on her own noisy signal z_{ijt} . Her estimate will be:²⁰

$$E[x_{jt}|z_{ijt}, \theta_t] = m_j z_{ijt} + (1 - m_j)E[x_{jt}|\theta_t].$$

If $m_j = 0$ (i.e., if the noise parameter $q_j \rightarrow \infty$), then the agent has no new information about x_{jt} and sticks to her prior $E[x_{jt}|\theta_t]$ when observing z_{ijt} . By contrast, if $m_j = 1$ (i.e., if the noise parameter $q_j = 0$), then the agent perfectly observes x_{jt} and ignores her own prior in her expectation of x_{jt} . In this sense, m_j captures how much information agents obtain about the j th component. We therefore call m_j the *attention* dedicated to the j th component.

While we have motivated our definition of m_j in the hypothetical case where agents condition on the latent factor θ_t , these quantities also determine agents' expectations about θ_t .

LEMMA 1: For each agent $i \in [0, 1]$, expectations about the latent factor θ_t satisfy

$$(11) \quad E_{it}[\theta_t] = E_{it-1}[\theta_t] + \sum_j g_j (z_{ijt} - E_{it-1}[z_{ijt}]),$$

where $g_j = \text{var}[\theta_t|\Omega_{it}](a_j/b_j^2)m_j$ denotes the weight placed on signal z_{ijt} .

The lemma confirms that attention coefficients m_j drive agents' responses to new information. The agent responds to each of her signals at date t in proportion to the Kalman gain g_j . This gain is the product of the steady state variance of θ_t and a measure of the precision of signal z_{ijt} , which is in turn proportional to attention m_j .²¹

²⁰We assume that all individuals choose the same attention allocation m_j . This is true in our model of optimal attention choice in Sections III and IV. It is also a standard assumption in the information choice literature (see, for example, Veldkamp 2011 and the references therein).

²¹To see why g_j captures the precision of z_{ijt} , consider the normalized signal $\hat{z}_{ijt} = z_{ijt}/a_j = \theta_t + \xi_{ijt}$, with $\xi_{ijt} = (b_j u_{jt} + q_j \epsilon_{ijt})/a_j$. The standard Gaussian updating formula implies that the gain on \hat{z}_{ijt} is proportional to the precision (inverse variance) of ξ_{ijt} . The proof of Lemma 1 shows that this precision equals $(a_j^2/b_j^2)m_j$.

C. Attention, Overreactions, and Underreactions

We now derive the coefficients for extrapolation in (1) and underreaction in (2) that an econometrician would estimate for this economy. The coefficient on current output in (1) satisfies

$$(12) \quad \gamma = \text{cov}[y_{t+k} - E_{it}y_{t+k}, y_t] \text{var}[y_t]^{-1} = d_0 \text{cov}[\theta_t - E_{it}\theta_t, y_t],$$

where $d_0 = (\rho^k \sum_j a_j) \text{var}[y_t]^{-1} > 0$, and $E_{it}y_{t+k} = f_{it}y_{t+k}$ denotes the k -period ahead forecast of output. Since agents are rational, their forecasts are equal to their conditional expectations. The equality in (12) follows because y_{t+k} depends only on θ_t and on shocks that are uncorrelated with date- t information. We note that the sign of γ is determined only by the covariance between the tracking error $\theta_t - E_{it}\theta_t$ and current output.

Meanwhile, the coefficient δ on the average forecast revision in (2) is

$$(13) \quad \begin{aligned} \delta &= \text{cov}[y_{t+k} - E_{it}y_{t+k}, \bar{E}_t y_{t+k} - \bar{E}_{t-1} y_{t+k}] \text{var}[\bar{E}_t y_{t+k} - \bar{E}_{t-1} y_{t+k}]^{-1} \\ &= d_1 \text{cov}[\theta_t - E_{it}\theta_t, \bar{E}_t \theta_t - \bar{E}_{t-1} \theta_t], \end{aligned}$$

where $d_1 = (\rho^k \sum_j a_j)^2 \text{var}[\bar{E}_t y_{t+k} - \bar{E}_{t-1} y_{t+k}]^{-1} > 0$. Hence, the sign of δ is determined only by the covariance between the tracking error of θ_t and the latest average forecast revision.

We start with two stark examples that demonstrate how the two covariances in (12) and (13) depend on individuals' attention choices. This, in turn, allows us to provide a simple illustration of the mechanisms behind our main results.

Example 1 (Asymmetric Attention and Extrapolation): Suppose that output has two components with $y_t = x_{1t} + x_{2t}$, and that the first component is procyclical with $a_1 > 0$. Agents pay full attention to the first component and none to the second ($m_1 = 1, m_2 = 0$). Then, the extrapolation coefficient in (12) becomes

$$\begin{aligned} \gamma &= d_0 \text{cov}[\theta_t - E_{it}\theta_t, x_{1t} + x_{2t}] \\ &= d_0 \text{cov}[\theta_t - E_{it}\theta_t, x_{2t}] = d_0 \text{cov}[\theta_t - E_{it}\theta_t, a_2 \theta_t] = a_2 d_0 \text{var}[\theta_t | \Omega_{it}], \end{aligned}$$

where the first equality follows from $\text{cov}[\theta_t - E_{it}\theta_t, x_{1t}] = 0$ for all agents $i \in [0, 1]$, because each agent is fully rational and observes x_{1t} perfectly. The second equality follows from $\text{cov}[\theta_t - E_{it}\theta_t, x_{2t}] = a_2 \text{cov}[\theta_t - E_{it}\theta_t, \theta_t]$, while the third one is due to individual rationality implying $\text{cov}[\theta_t - E_{it}\theta_t, \theta_t] = \text{cov}[\theta_t - E_{it}\theta_t, \theta_t - E_{it}\theta_t]$. We conclude that $\gamma = a_2 d_0 \text{var}[\theta_t | \Omega_{it}]$, and thus that the extrapolation coefficient γ has the same sign as a_2 .

Example 1 shows that the econometrician will find extrapolation, i.e., overreactions to current output ($\gamma < 0$), if and only if $a_2 < 0$; that is, if and only if the component x_{2t} , to which agents pay no attention, is *countercyclical*. This highlights how our rational model can generate overreactions. In effect, the example shows that

the *overreaction* to recent output documented in the survey data can be interpreted as an *underreaction to countercyclical components*.

The economic intuition behind this fact, which captures one of the main ideas of this paper, is as follows: When output y_t is high, the procyclical component x_{1t} , all else equal, also tends to be high, which represents good news about the latent factor θ_t . However, the countercyclical component x_{2t} , on average, also tends to be large, which dampens any good news about the latent factor. When agents pay relatively less attention to countercyclical components, their posteriors place only a small weight on this dampening effect. As a result, when output is high, agents tend to be more optimistic than the econometrician (who controls for total output) about the future. This leads to a seeming extrapolation, which manifests itself in a negative correlation between future forecast errors and current output.

Our second example shows that our environment, despite such overreactions, remains consistent with the underreactions documented in Section I.

Example 2 (Limited Attention and Underreactions): Consider the setting in Example 1, but now suppose that agents' attention to the first component of output is also limited: $0 < m_1 < 1$. Since the average revision is $\bar{E}_t \theta_t - \bar{E}_{t-1} \theta_t = \int_0^1 (E_{jt} \theta_t - E_{jt-1} \theta_t) dj$, the linearity of the covariance operator and the symmetry of attention choices imply that

$$\begin{aligned} \delta &= d_1 \text{cov}[\theta_t - E_{it} \theta_t, \bar{E}_t \theta_t - \bar{E}_{t-1} \theta_t] \\ &= d_1 \text{cov}[\theta_t - E_{it} \theta_t, E_{jt} \theta_t - E_{jt-1} \theta_t] = d_1 \text{cov}[E_{jt} \theta_t - E_{it} \theta_t, E_{jt} \theta_t - E_{jt-1} \theta_t], \end{aligned}$$

where the third equality follows by adding and subtracting agent j 's forecast error $\theta_t - E_{jt} \theta_t$, and noting that it is uncorrelated with j 's forecast revision. We conclude that $\delta > 0$ if, for all i and $j \neq i$, $\text{cov}[E_{jt} \theta_t, E_{jt} \theta_t - E_{jt-1} \theta_t] > \text{cov}[E_{it} \theta_t, E_{jt} \theta_t - E_{jt-1} \theta_t]$. This always holds in our example. Intuitively, when $m_1 < 1$, agent i and j observe *different* signals, which makes agent j 's forecast revision more strongly correlated with her own expectation.

This second example shows that the econometrician will estimate underreactions to average forecast revisions ($\delta > 0$) when agents' attention to at least one component is limited. This extends the results in Coibion and Gorodnichenko (2015) to our case.²² The intuition is as discussed above. As long as information is dispersed, rational individuals respond less strongly to *average new information* than agents in the fully informed rational benchmark. This leads to underreactions of expectations similar to those documented in the survey data.

Combined, the examples above demonstrate how attention choices map into the over- and underreaction coefficients γ and δ , respectively. Specifically, they show how limited, asymmetric attention to a procyclical component can explain the

²²The baseline model in Coibion and Gorodnichenko (2015) assumes uncorrelated noise terms across agents. In an extension, Coibion and Gorodnichenko (2015, online Appendix A) notes that the coefficient δ measured by an econometrician will be attenuated by the presence of common noise terms u_{jt} . A novel result in this example and Proposition 2 that follows is that, despite this effect, we always have $\delta > 0$.

simultaneous over- and underreactions of survey expectations ($\gamma < 0$ and $\delta > 0$). Using similar steps, Proposition 2 extends our results to the general case with N components and arbitrary attention choices.

PROPOSITION 2: *Output forecasts overreact to current output ($\gamma < 0$ in (1)) if and only if agents pay asymmetric attention to procyclical components, so that $\sum_j a_j(1 - m_j) < 0$. Output forecasts underreact to new information on average ($\delta > 0$ in (2)) if and only if attention is limited, i.e., if there exists $j \in \{1, \dots, N\}$ such that $0 < m_j < 1$.*

The first part of the proposition states the key sufficient statistic: $\sum_j a_j(1 - m_j)$. Our model is consistent with overreactions to current output (i.e., extrapolation) whenever this statistic is negative. This is clearly the case when agents are inattentive ($m_j \simeq 0$) to components that are countercyclical, which covary negatively with the latent factor ($a_j < 0$), and are more attentive to procyclical components ($a_j > 0$). Thus, asymmetric attention to procyclical components is a *sufficient* condition for extrapolation ($\gamma < 0$).

The proposition further implies that asymmetric attention is also a *necessary* condition for extrapolation. If attention were symmetric with $m_j \equiv \bar{m}$ for all j , then we would have $\sum_j a_j(1 - \bar{m}) \geq 0$, since $\sum_j a_j > 0$, and hence $\gamma \geq 0$. Intuitively, the symmetric case is similar to the rational benchmark with noisy information studied in Section IB, where rational updating induces underreactions in both (1) and (2). Hence, the symmetric case is inconsistent with the large body of evidence documenting extrapolation.

The second part of the proposition extends the results of Coibion and Gorodnichenko (2015) to our framework. We find that underreactions to new information occur whenever attention is limited for at least one component.

D. Summary and Extensions

In summary, our model is able to match the stylized facts whenever attention is both *limited* and *asymmetric*. We close this section by discussing two important extensions.

First, we have presented a latent factor model with several components of output. This classical structure conveys our main contribution and leads naturally to our macroeconomic example in Section IV. However, the model in this section is not the only possible parametrization in which asymmetric attention explains the patterns that we find in the data. In particular, Proposition B.1 in online Appendix B fully characterizes the coefficients in (1) and (2) for a larger class of linear models, in which we allow for (i) the direct effects of several, latent factors on output; (ii) the correlation between component-specific shocks; and (iii) the explicit observation of (and dependence on) lagged outcomes. This extension, which encompasses most linear macroeconomic models, delivers necessary and sufficient conditions for over- and underreactions based on limited, asymmetric attention more generally.

Second, we have focused our discussion of forecast revisions on (2), which is the regression of forecast errors on *average* forecast revisions proposed by Coibion and Gorodnichenko (2015). By contrast, in contemporaneous and closely related

work, Bordalo et al. (2020) considers the regression of forecast errors on *individual* forecast revisions:²³

$$(14) \quad y_{t+k} - f_{it}y_{t+k} = \alpha + \delta^{ind}(f_{it}y_{t+k} - f_{it-1}y_{t+k}) + \xi_{it}.$$

Using a range of survey data, Bordalo et al. (2020) estimates that $\delta^{ind} < 0$, which is inconsistent with the predictions of our baseline model, and also with other models with rational expectations in which agents recall their own forecast revisions. Table C.1 in the online Appendix reports estimates of (14) for output forecasts in the US SPF. We estimate overreactions to individual revisions ($\delta^{ind} < 0$), but unlike our estimates of (1) and (2), which motivate our analysis, this result appears sensitive to outliers.²⁴ Online Appendix F considers an extension of our framework, which allows for both asymmetric attention and irrational overconfidence (e.g., Moore and Healy 2008 and Broer and Kohlhas 2019). We show that, when one introduces a small bias, the extended model can account not only for the stylized facts that we have emphasized ($\gamma < 0$ in (1) and $\delta > 0$ in (2)), but also for overreactions to individual revisions ($\delta^{ind} < 0$ in (14)). Crucially, the extended model can fit these empirical patterns *only if* one introduces asymmetric attention. As discussed in Section IB, the baseline model in Bordalo, Gennaioli, and Shleifer (2018) predicts that γ and δ have the same sign. Thus, regardless of whether there are overreactions to individual revisions, asymmetric attention is *necessary* to reconcile the varied survey evidence within the class of models examined.

So far, we have considered reduced-form economies. In deriving our results, we have taken agents' attention choices, as summarized by the set of m_j , as given. We now move on to studying the potential sources of asymmetric attention.

III. Attention Choices

In this section, we consider agents' attention choices. We show that attention gravitates towards volatile components that are important to decision-makers. Combined with our previous results, this demonstrates that a rational theory of limited attention can match the survey evidence when procyclical components are either more volatile or more important.

A. A Model with Attention Choice

We augment our environment to incorporate attention choice. To do so, we assume the following timing of events: In an initial period $t = 0$, each agent first chooses her attention allocation m_j to the different components x_{jt} of output (or equivalently, the noise terms q_j). She makes this choice *ex ante*, behind the veil of ignorance. The agent then receives a (infinitely) long sequence of signals, denoted by z_t^0 . This assumption ensures that the agent's signal extraction problem is initialized in steady

²³ See also Fuhrer (2017) and Broer and Kohlhas (2019) for related results using inflation forecasts.

²⁴ Indeed, we cannot reject that $\delta^{ind} = 0$ once we remove outliers in the top 1 percent of forecast errors and revisions. This is in contrast to our estimates of (1) and (2). See also Angeletos and Huo (2021) for similar empirical results using inflation forecasts.

state. In each subsequent period $t > 0$, agent i 's information set is the history of her past signals $\Omega_{it} = \{z_i^0 \cup (z_{i1s}, \dots, z_{iNs})_{s=1}^{s=t}\}$. After observing the latest signal vector z_{it} in each period t , the agent chooses an action a_{it} .

The agent's lifetime utility is

$$(15) \quad E_{i0} \sum_{t=1}^{\infty} \beta^t \mathcal{U}_{it}, \quad \mathcal{U}_{it} = -(a_i^* - a_{it})^2 - K(m),$$

where $\beta \in (0, 1)$ denotes the time discount factor. The agent's per-period utility \mathcal{U}_{it} , in (15) consists of two terms. The first term is a quadratic loss that the individual incurs when she deviates from her ideal action a_i^* . The second term reflects the cost of attention $K(m)$. We assume that $K(\cdot)$ is positive, increasing in all m_j , and convex. We further assume that the ideal action, which the agent would take under full information about all stochastic disturbances, can depend both on the unobserved latent factor and on the structural components:

$$(16) \quad a_i^* = w_\theta \theta_i + \sum w_{xj} x_{ji},$$

where $w_\theta \in \mathbb{R}$ and $w_{xj} \in \mathbb{R}$ for all j . With these preferences, the optimal choice of an agent who has information Ω_{it} in the last stage at date t is to set $a_{it} = E[a_i^* | \Omega_{it}]$.

Equations (15) and (16) nest the benchmark case in which agents care only about forecasting future output as accurately as possible: when $w_\theta = \rho^k \sum_j a_j$ and $w_{xj} = 0$, a_i^* becomes the full-information mean squared optimal forecast of y_{t+k} , which is $E_t^{FIRE}[y_{t+k}] = \rho^k \sum_j a_j \theta_t$. However, (15) and (16) also allow us to capture more general cases in which agents' ideal choice depends differently on the various structural components of output. This allows us to account for cases in which agents do not necessarily design their attention choices with the objective of predicting future output as accurately as possible. Instead, agents can also skew their attention choices towards the components of output that are the most important for their own specific decision problems. A firm, for example, might choose to pay more attention to its own sector than the economy as a whole (see Section IV for a related example).

B. Optimal Attention to Important and Volatile Variables

We now derive agents' attention choices. To do so, it is instructive to first derive agents' ex ante expected utility as a function of their attention choices.

LEMMA 2: *Each agent's ex ante expected lifetime utility in the initial period $t = 0$ equals*

$$(17) \quad \frac{1}{1-\beta} E[\mathcal{U}_{it}] \propto -V[a_i^* | \Omega_{it}] - K(m)$$

$$(18) \quad = -\sum_j w_{xj}^2 b_j^2 (1 - m_j) - \text{var}_t[\theta_t] \left[w_\theta + \sum_j w_{xj} a_j (1 - m_j) \right]^2 - K(m).$$

Lemma 2 first provides a natural characterization of an agent's ex ante expected utility (i.e., before she observes her signals z_i^t). Intuitively, for every realization of her signals at date t , the agent will set $a_{it} = E[a_i^* | \Omega_{it}]$. Hence, her maximized utility

depends on the expected squared deviation of $E[a_t^*|\Omega_{it}]$ from a_t^* , which reduces to the conditional variance in (17). Lemma 2 then derives an expression for the conditional variance, using the law of total variance:

$$\text{var}[a_t^*|\Omega_{it}] = \text{var}[a_t^*|\Omega_{it}, \theta_t] + \text{var}[E[a_t^*|\Omega_{it}, \theta_t]|\Omega_{it}].$$

Accordingly, the first term in (18) reflects the uncertainty about the optimal action conditional on the latent factor. It equals the sum of the conditional variances $\text{var}[x_{jt}|\Omega_{it}, \theta_t]$ across the components x_{jt} , weighted by their importance w_{xj} in agents' utility. The uncertainty about each component naturally increases in its volatility b_j^2 but decreases in agents' attention m_j .

The second term in (18) measures the residual uncertainty $\text{var}[\theta_t|\Omega_{it}] \equiv \text{var}_t[\theta_t]$, scaled by the uncertainty about the ideal action $a_t^* = w_\theta \theta_t + \sum_j w_{xj} x_{jt}$ that is attributable to θ_t (i.e., by the term in square brackets). We provide a brief derivation of $\text{var}_t[\theta_t]$, to show how it depends on agents' attention choices. In turn, combined with (16) and (18), this will then allow us to derive an expression for agents' optimal attention choices.

Recall that the effective precision of signal z_{ijt} about θ_t is $\tau_j = a_j^2/(b_j^2 + q_j^2)$, and let

$$(19) \quad \tau(m) = \sum_j \tau_j$$

denote the total precision of date t signals. Starting at date t , the conditional variance about next period's fundamental is $\text{var}_t[\theta_{t+1}] = \rho^2 \text{var}_t[\theta_t] + \sigma_\theta^2$. After updating based on date $t+1$ signals, this variance satisfies the linear precision rule $\text{var}_{t+1}[\theta_{t+1}]^{-1} = \text{var}_t[\theta_{t+1}]^{-1} + \tau(m)$. Solving for a steady state where $\text{var}_t[\theta_t] = \text{var}_{t+1}[\theta_{t+1}] = V$ then delivers

$$\sigma_\theta^2 = V[1 - \rho^2 + \tau(m)\sigma_\theta^2] + V^2\tau(m)\rho^2.$$

Thus, the total precision τ of an agent's signals is a sufficient statistic for her uncertainty about the latent factor, and we can write

$$(20) \quad \text{var}_t[\theta_t] = V[\tau(m)],$$

where $V'(\tau) < 0$ and $\partial\tau/\partial m_j > 0$ from (19). Combined, (18) and (20) allow us to characterize agents' attention choices. Proposition 3 summarizes the results.

PROPOSITION 3: *Agents' optimal attention choices satisfy, for all j such that $0 < m_j < 1$,*

$$(21) \quad w_{xj}^2 b_j^2 + \mu_\tau a_j^2 b_j^{-2} + \mu_\alpha w_{xj} a_j = \frac{\partial K(m)}{\partial m_j},$$

where $\mu_\tau > 0$ and $\mu_\alpha > 0$ denote Lagrange multipliers.

Proposition 3 uses the fact that optimal (interior) attention choices equate the marginal benefit of paying more attention to each component to its marginal cost.

The marginal benefit on the left-hand side of (21) consists of three terms. The first term is the benefit of resolving uncertainty about the optimal action conditional on θ_t . This benefit is higher for components that are more important for the optimal action (high w_j) and more volatile (high b_j).

The second and third terms capture a more nuanced effect: by learning about x_{jt} , the agent also acquires information about the latent factor θ_t , which generates *learning spillovers* by resolving uncertainty about x_{kt} for $k \neq j$. The second term measures the effect of attention m_j on the effective precision τ of agents' signals about θ_t . The multiplier μ_τ is the shadow value of increasing this precision. This benefit of attention is larger for components that are highly correlated with the fundamental (high a_j^2), but spillovers are attenuated for components that are highly volatile (high b_j^2). The third term measures an adjustment to this effect, namely, that information about θ_t becomes less valuable to an agent if she already has precise information about the structural components x_{jt} , and hence about her optimal action. The multiplier μ_α is the shadow value of reducing the residual uncertainty about a_t^* that is attributable to θ_t .

While these effects are subtle, the underlying intuition is clear. On one hand, agents are more likely to pay attention to components that are important for their utility, those with large weights w_{xj} in (16). On the other hand, agents also prefer to pay attention to volatile components (with a high idiosyncratic variance b_j^2), as long as learning spillovers are not too strong. This tendency for attention to gravitate towards important and volatile variables is familiar from much of the literature on information choice (Veldkamp 2011), and has recently received additional empirical support in microlevel firm data (Coibion, Gorodnichenko, and Kumar 2018b). Proposition 3 confirms that this intuition carries over to our component-based model.

Figure 4 provides a numerical example, which illustrates the effects of component volatility and utility weighting on agents' optimal attention choices. To demonstrate the role of learning spillovers, the figure considers three scenarios for the variance σ_θ^2 of the latent factor. Intuitively, spillovers are minimized when the variance of the latent factor θ_t is small. The two panels confirm the main points in our discussion: The relative attention m_1^*/m_2^* paid to component 1 increases as this component becomes more volatile ($\uparrow b_1$ in panel A) and more important in agents' objective function ($\uparrow w_{x1}$ in panel B). In both cases, the rate of increase is smaller when there are strong spillovers (high σ_θ^2). This reflects the intuition that strong learning spillovers incentivize an agent to push on all margins to learn more about the latent factor, which in turn leads her to respond less strongly to component-specific features.

We have so far kept the functional form of the attention cost function $K(m)$ general. Online Appendix G derives the first-order condition (21) explicitly for an entropy-based cost function, and shows that the main comparative statics remain the same. In addition, we show that an entropy-based cost function naturally yields limited attention choices $m_j < 1$, because it implies that the marginal cost of full attention is infinite ($\lim_{m_j \rightarrow 1} \partial K(m)/\partial m_j = \infty$).

In sum, asymmetric attention arises naturally from costly attention choice if some components are either more volatile, or more important to decision-makers. Combined with the insights of the previous section, we can therefore conclude that a

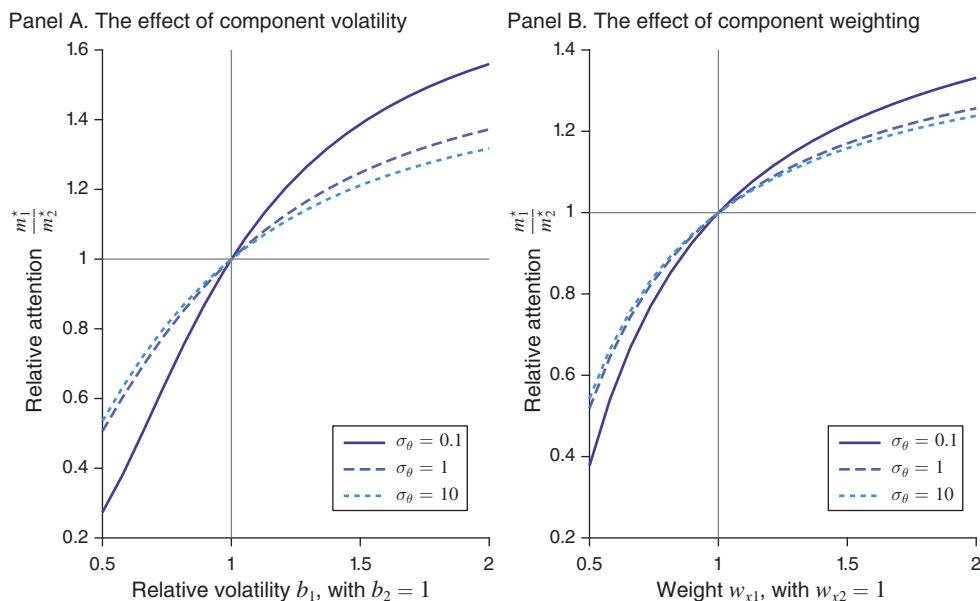


FIGURE 4. OPTIMAL ATTENTION: NUMERICAL EXAMPLE

Notes: The charts show the properties of optimal attention choices as a function of component volatilities b_j , utility weighting w_{xj} , and the variance σ_θ^2 of the latent factor in a numerical example with two components. The parameters not detailed in the figure are set at $a_1 = a_2 = 1$, $\rho = 0.9$, $w_\theta = 0$. The cost function $K(m)$ is set to the reduction in entropy, as derived in Proposition G.1 in the online Appendix.

rational theory of limited attention can match the survey evidence when procyclical components are either more volatile or more important. In the next section, we apply this reasoning to a simple macroeconomic model and show that, for reasonable parameters, attention gravitates to procyclical variables.

Before moving on to the application, we consider two more points. First, we explore an alternative model of information choice in which agents have full flexibility in their information design. Second, we revisit the data and show that the survey evidence is consistent with an additional prediction of our framework.

C. Fully Flexible Information Choice

Proposition 3 characterizes the solution to a *constrained* information choice problem. Equation (9) restricts agents to acquire N separate, conditionally independent signals z_{ijt} about the components x_{jt} of output. This is one of two popular approaches. An alternative approach is to instead allow agents full flexibility when designing the conditional distribution of their signals given the state of the economy (e.g., Sims 2003). The choice between the two approaches is typically made based on the problem at hand, and on tractability. In the context of our analysis, it is interesting to compare the predictions of each approach.

Exploiting the characterizations of optimal signal design in Maćkowiak, Matějka, and Wiederholt (2018), online Appendix H shows that agents in our model, when

equipped with an entropy-based cost function, would optimally choose to receive a single signal of the optimal action:²⁵

$$(22) \quad s_{it}^* = a_t^* + h\xi_t + q^*\epsilon_{it},$$

where h depends on the utility weights w_θ and w_{xj} , ξ_t is a Gaussian white noise sequence that depends on the common shocks η_t and u_{jt} , and q^* denotes a scalar that depends on the cost of attention $K(m)$. Equation (22) shows that the asymmetry of attention now depends on the weight w_{xj} in agents' optimal action both through their influence on a_t^* and the coefficient h in the optimal signal. This has important empirical implications.

For example, consider the benchmark case in which agents' utility in (15) is equivalent to the mean-squared error of next period's output forecast ($w_\theta = \rho \sum_j a_j$ and $w_{xj} = 0$, as discussed above). In this case, it follows that $h = 0$ in (22).²⁶ As a result, the fully optimal signal boils down to $s_{it}^* = (\rho \sum_j a_j) \theta_t + q^* \epsilon_{it}$, which is a simple noisy signal of θ_t . Similar to the results in Proposition 2, and due to the symmetry of underlying preferences, such a signal is inconsistent with extrapolation.²⁷ Indeed, in this case, agents systematically under-react to new information about current output, yielding $\gamma > 0$ in (1).

Consider now instead the case in which the weights w_{xj} in agents' optimal action are asymmetric across the structural components. In this case, agents' forecasts of future output given s_{it}^* can exhibit extrapolation. Similar to the results in Proposition 2, this occurs when the weights w_{xj} are tilted towards procyclical components. This is easiest to see in the following example, which extends our previous Example 1 to flexible information choice.

Example 3 (Asymmetric Attention and Extrapolation, cont.): As in Example 1, suppose that output has two components, where $a_1 > 0$ and $a_2 < 0$. Agents' ideal action depends only on the first component ($w_{x1} > 0$ while $w_{x2} = w_\theta = 0$). Building on the results in Maćkowiak, Matějka, and Wiederholt (2018),

²⁵Heuristically, one can apply the results in Maćkowiak, Matějka, and Wiederholt (2018) after expressing a_t^* as an ARMA process in reduced form. In particular, substituting (7) and (8) into (16) shows that

$$\begin{aligned} a_t^* &= \underbrace{(w_\theta + \sum w_{xj} a_j)}_{\equiv \bar{w}_\theta} \theta_t + \underbrace{\sum w_{xj} b_j}_{\equiv \bar{w}_{xj}} u_{jt} = \rho a_{t-1}^* + \bar{w}_\theta \eta_t + \bar{w}_x' u_t - \rho \bar{w}_x' u_{t-1} \\ &\equiv \rho a_{t-1}^* + c_0' v_t + c_1' v_{t-1}. \end{aligned}$$

Hence, a_t^* is an autoregressive-moving-average model (ARMA) process with a vector of innovations $v_t = [\eta_t, u_t]'$. In online Appendix H, we further demonstrate that this process can be represented as a standard ARMA process with scalar innovations ξ_t , so that we can apply the characterization of optimal signals provided in Maćkowiak, Matějka, and Wiederholt (2018).

²⁶See Proposition H.1 in the online Appendix, or Cover and Thomas (2012) for the standard result in which the optimal action a_t^* is proportional to a simple AR(1) process.

²⁷Consider the extrapolation coefficient in (1) based on $s_{it}^* = (\rho \sum_j a_j) \theta_t + q^* \epsilon_{it}$. It follows that

$$\begin{aligned} \gamma &= \text{cov}(y_{t+k} - E_{it} y_{t+k}, y_t) \text{var}[y_t]^{-1} \\ &= d_0 \text{cov}(\theta_t - E_{it} \theta, y_t) = d_0 \sum_j a_j \text{var}[\theta_t | s_i^{*,t}] > 0, \end{aligned}$$

where we have also used that $y_t = \sum_j a_j \theta_t + \sum_j b_j u_{jt}$ and that $\sum_j a_j > 0$.

Proposition H.1 in the online Appendix shows that, if the costs of attention are sufficiently small, the optimal signal tends to $s_{it}^* = x_{1t} + q^* \epsilon_{it}$. Hence, the information structure is identical to that in Examples 1 and 2, where $0 < m_1 < 1$ and $m_2 = 0$. The arguments in Example 1 and 2 now imply that $\gamma < 0$ and $\delta > 0$. By continuity, the model with flexible information choice generates $\gamma < 0$ and $\delta > 0$ as long as the weight w_{x1} is sufficiently large relative to w_θ and w_{x2} .

Combined, these examples show that we cannot test, based on survey data alone, whether the asymmetry of attention is driven by conditionally independent signals or by a flexibly designed, skewed signal. We only know that the fully flexible case is rejected by the data if agents care exclusively about the mean-squared error of output forecasts. By contrast, Proposition 3 shows that the “conditionally independent signals” structure, even in the mean-squared error case, can be consistent with the simultaneous over- and underreactions documented in the data, so long as there are differences in the volatility of the underlying components.

D. Are Attention Choices Optimal? Supplementary Evidence

We briefly return to the data to compare the quality of agents’ expectations to that of standard time-series models. Figure 5 shows updated values from Stark (2010), available from the Federal Reserve Bank of Philadelphia’s website.²⁸ The chart illustrates the *relative root mean-squared error* (RRMSE) of one-quarter- and four-quarter-ahead forecasts of output growth from US SPF relative to three optimally chosen time-series models. An RRMSE ratio below unity indicates that the SPF consensus forecast is more accurate. All time-series models fall short of survey forecasts at the one-quarter horizon, while the more sophisticated ARMA models achieve a close match with the SPF at the four-quarter horizon.

This supplementary evidence suggests that forecasters do better than simple time-series models at forecasting output. This is consistent with our model, in which agents pay attention to underlying, structural components of the forecasted variable, but inconsistent with a model where agents consider only the past time series of output (see, for instance, Proposition 11.2 in Lütkepohl 2007). In addition, this evidence rejects a simple behavioral story where agents derive forecasts from a misspecified ARMA model. Recent behavioral theory, such as Bordalo et al. (2020), is more nuanced, and further work would be needed to test whether forecasts in the data are more or less accurate than such theories predict. Hence, we interpret the supplementary evidence as a sanity check, which implies that our theory is consistent with moments of the data beyond the motivating evidence in Section I.

We now turn to an application of our ideas to a standard macroeconomic model.

IV. A Macroeconomic Example

In this section, we illustrate the sources and effects of asymmetric attention in a flexible-price business cycle model. We analyze an environment in which firms

²⁸<https://www.philadelphiafed.org/research-and-data/real-time-center.html>.

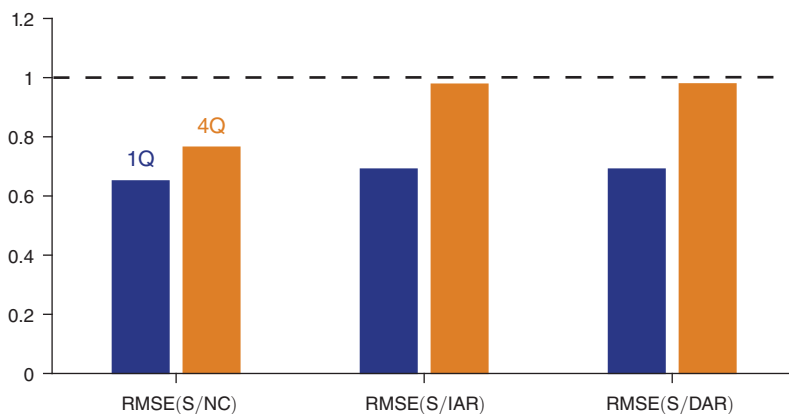


FIGURE 5. FORECAST PRECISION RELATIVE TO TIME-SERIES MODELS

Notes: The chart shows updated values from Stark (2010), available from the Federal Reserve Bank of Philadelphia's website. The chart illustrates the *relative root mean-squared error* (RRMSE) of one-quarter- and four-quarter-ahead forecasts of output growth from the US SPF (S) relative to three time-series models: NC denotes a random walk forecast, IAR forecasts from an ARMA model chosen to minimize one-quarter-ahead forecast errors, and DAR forecasts from ARMA models chosen to minimize forecast errors at each forecast horizon. The sample period is 1985:I–2015:II. An RRMSE ratio below unity indicates that the SPF consensus forecast is more accurate.

choose output under imperfect information. We show that firms' output choices can be decomposed into two components: First, a *productivity component*, which summarizes the effects of a firm's own productivity; and second, an *aggregate supply component*, which captures the effects of other agents' behavior on an individual firm's output choice. We document that, for standard parameters, the productivity component is procyclical, while the aggregate supply component is countercyclical. In accordance with the evidence in Coibion, Gorodnichenko, and Kumar (2018b), we show that firms' attention choices are asymmetric and tend to abstract from the aggregate component. As a result, and in line with the analysis above, we find that firms' expectations of output mirror the estimated extrapolation and underreactions from the survey data. Finally, we show that asymmetric attention leads to more volatility and persistence in output.

A. Model Setup

The economy consists of a representative household and a continuum of monopolistically competitive firms $i \in [0, 1]$, which specialize in the production of differentiated goods.

Households.—The representative household has lifetime utility

$$(23) \quad E_0 \sum_{t=0}^{\infty} \beta^t [\log C_t - \xi_t N_t], \quad \xi_t > 0,$$

where β denotes the time discount factor, C_t the consumption index at time t , N_t the number of hours worked by the household, and ξ_t a shock to the disutility of

labor. The consumption index C_t and associated welfare-based price index P_t are

$$(24) \quad C_t = \left[\int_0^1 C_{it}^{\frac{\sigma-1}{\sigma}} di \right]^{\frac{\sigma}{\sigma-1}}, \quad P_t = \left[\int_0^1 P_{it}^{1/(\sigma-1)} di \right]^{\sigma-1},$$

where C_{it} is the amount the household consumes of goods produced by firm i at price P_{it} , and $\sigma > 1$. The household's per-period budget constraint is

$$(25) \quad \int_0^1 P_{it} C_{it} di + B_{t+1} \leq \int_0^1 \Pi_{it} di + W_t N_t + (1 + R_t) B_{t+1} + T_t^h,$$

where Π_{it} denotes the profits of firm i , W_t the nominal wage, R_t the nominal rate of return on riskless bonds, B_t its holdings of riskless bonds, and T_t^h lump-sum nominal transfers. The representative household's objective is to maximize its utility (23) subject to (25).

Firms.—A representative firm $i \in [0, 1]$ chooses its output Y_{it} to maximize its own expectation of the household's valuation of its profits, using the stochastic discount factor $(P_t C_t)^{-1}$. The expected valuation of profits at time t is equal to

$$(26) \quad \mathcal{V}_{it} = E_{it} \left[\frac{1}{P_t C_t} \Pi_{it} \right], \quad \Pi_{it} = P_{it} Y_{it} - W_t N_{it},$$

where the inverse-demand for a firm's product is consistent with household optimality: $P_{it} = P_t (Y_{it}/Y_t)^{-1/\sigma}$. Firm output is produced in accordance with the production function

$$(27) \quad Y_{it} = A_{it} N_{it}^\alpha, \quad \alpha \in (0, 1),$$

where N_{it} denotes the amount of labor input used and A_{it} firm-specific productivity.

Shocks.—We let lowercase letters denote natural logarithms of their uppercase counterparts. Firm-specific productivity $a_{it} = \log A_{it}$ is

$$(28) \quad a_{it} = \theta_t + u_t^x + \epsilon_{it}^a,$$

where the persistent, common component θ_t follows an AR(1) process,

$$(29) \quad \theta_t = \rho \theta_{t-1} + u_t^\theta, \quad u_t^\theta \sim \mathcal{N}(0, \sigma_\theta^2),$$

while the transitory and firm-specific components are distributed as $u_t^x \sim \mathcal{N}(0, \sigma_x^2)$ and $\epsilon_{it}^a \sim \mathcal{N}(0, \sigma_a^2)$, respectively. This is similar to the decomposition used in Kydland and Prescott (1982). The household's disutility of labor is subject to a transitory shock with

$$(30) \quad \log \xi_t = \bar{\xi} + u_t^n, \quad u_t^n \sim \mathcal{N}(0, \sigma_n^2),$$

where $\bar{\xi} \in \mathbb{R}$. We show below that the labor supply shock introduces a component-specific innovation to aggregate output. In effect, u_t^n will play the role of

one of the component-specific disturbances u_{jt} discussed in Section II. We assume that the innovations u_t^x , u_t^θ , u_t^n , and ϵ_{it}^a are independent of each other, across time, and across firms.

Timeline.—In an initial period $t = 0$, firms choose ex ante how much attention to devote to the various components of output, which we define below. In each subsequent period $t > 0$, nature determines the realization of the shocks u_t^x , u_t^θ , u_t^n , and ϵ_{it}^a . The economy then proceeds through three stages. In the first stage, firms commit to their output choices. After output choices are sunk, the economy transitions to the second stage, in which the labor market opens. Each firm hires the amount of labor $n_{it} = \alpha^{-1}(y_{it} - a_{it})$ that is necessary to implement its previous output choice y_{it} conditional on realized productivity a_{it} .²⁹ The household observes its marginal disutility $\xi_t = \bar{\xi} + u_t^n$ of labor and the permanent productivity component θ_t , and then makes its labor supply choice.³⁰ The real wage adjusts to clear the labor market. In the third and final stage, goods markets open, goods prices adjust, and the household consumes.

Information Structure.—To complete the description of the economy, it is necessary to specify the information structure and firms' associated attention choice problem. Our assumptions are based on the following decomposition of firms' expected profits.

PROPOSITION 4: *A second-order approximation of firm i 's expected discounted profits satisfies*

$$(31) \quad v_{it} \propto -\frac{1}{2} E_{it}[(y_{it} - y_{it}^*)^2],$$

where the firm's ideal output under full information y_{it}^* can be decomposed into

$$(32) \quad y_{it}^* = x_{i1t} + x_{2t}$$

with

$$(33) \quad x_{i1t} = ra_{it}, \quad x_{2t} = \alpha r(\sigma^{-1}y_t - \omega_t),$$

and where ω_t denotes the real wage, $y_t = \int_0^1 y_{it} di$, and $r \equiv \sigma/(\sigma + \alpha(1 - \sigma)) > 1$.

²⁹ We assume that firms do not update beliefs a second time after observing their labor input. That is, we assume that firms do not invert their production function to back out a second signal of a_{it} . This assumption is common in other models of attention choice, where firms' actions are pre-set for the period (e.g., Maćkowiak and Wiederholt 2009). Vives and Yang (2018) provide a more detailed discussion of the need for such assumptions in models of costly attention choice to maintain imperfect information.

³⁰ Because the household does not observe the realization of u_t^x in the second stage, output will respond differently to innovations in θ_t and u_t^x . This friction creates a meaningful distinction between these two shocks. Without this friction, only shocks to the sum $\int_0^1 a_{it} di = \theta_t + u_t^x$ would matter for output. An equivalent way to create distinct dynamics would be to study a model in which one of the factors of production, such as capital, is predetermined before the realization of some of the shocks (see, for example, Angeletos, Iovino, and La'O 2016).

In the spirit of Lucas (1977) and Maćkowiak and Wiederholt (2009), equations (32) and (33) decompose each firm's ideal output choice into two components: We refer to x_{i1t} as the *productivity component*, since it depends on a firm's own productivity a_{it} . Clearly, each firm produces more when it is more productive. We refer to the second component, x_{2t} , as the *aggregate supply component*, which encapsulates the general equilibrium effects of other agents' behavior on an individual firm's output choice. The aggregate supply component, in turn, is comprised of two terms: On one hand, firms produce more when aggregate demand in the economy y_t is high. On the other hand, a firm also chooses to produce less when the real wage it faces ω_t is high. Both effects are captured in (33).³¹

Given this decomposition, our assumptions about firms' information sets and attention choices mirror those in our baseline model. Specifically, we assume that firm i 's information set consists of the history of component-based signals:

$$(34) \quad \Omega_{it} = \{z_i^0 \cup (z_{i1s}, z_{i2s})_{s=1}^{s=t}\},$$

where

$$(35) \quad z_{i1t} = x_{i1t} + q_1 \epsilon_{i1t}, \quad z_{i2t} = x_{2t} + q_2 \epsilon_{i2t},$$

and $\epsilon_{ijt} \sim \mathcal{N}(0, 1)$ is independently distributed across time and firms for $j = \{1, 2\}$. We also assume that in the initial period, after determining their attention choices, firms receive a (infinitely) long sequence of signals generated by (35), denoted by z_i^0 . This assumption once more ensures that the firms' signal extraction problem is initialized in steady state. Finally, as in our reduced-form framework, we assume that firms choose normalized attention parameters $m_j = \text{var}(x_{jt}|\theta_t) / (\text{var}(x_{jt}|\theta_t) + q_j^2)$ at a cost $K(m)$.

B. Equilibrium Characterization

We now proceed to characterize equilibrium output in the economy.

Equilibrium with Full Attention.—We start with the case in which firms pay full attention to both components (i.e., $m_j = 1$ for $j = 1, 2$) and there are no firm-specific productivity shocks ($\sigma_a = 0$). This special case illustrates some important findings, which will carry over to our numerical solution of the full model with limited attention. In this special case, Proposition 4 directly implies that each firm sets $y_{it} = y_{it}^* = x_{i1t} + x_{2t}$, so that

$$(36) \quad y_t = \int_0^1 y_{it} di = x_{1t} + x_{2t},$$

³¹ Unlike the similar decomposition used in Maćkowiak and Wiederholt (2009), the two components x_{i1t} and x_{2t} are correlated in this application. For example, a shock to θ_t will affect both components. Furthermore, in contrast to the baseline model from Section II, the error terms in the two components are also correlated, since both depend on the transitory productivity shock u_t^* . Hence, in order to characterize the properties of firms' expectations, we will use the more general results listed in Proposition B.1 in online Appendix B.

with $x_{1t} = \int_0^1 x_{i1t} di$. Thus, output has the same component-based structure as in the baseline model from Section II. The components x_{jt} of output can now further be characterized directly from (33). As for the productivity component x_{1t} , we have

$$(37) \quad x_{1t} = r\theta_t + ru_t^x.$$

This component is *procyclical*, since it places a positive weight $r > 0$ on the latent factor θ_t . Turning to the aggregate supply component x_{2t} , the real wage in equilibrium is $\omega_t = E_{ht}y_t + u_t^n$, where $E_{ht}[\cdot]$ denotes household expectations. Thus, we conclude from (33) and (36) that

$$(38) \quad \begin{aligned} x_{2t} &= \alpha r \left(\frac{1-\sigma}{\sigma} y_t + \left(\frac{1}{1-\alpha} - r \right) u_t^x - u_t^n \right) \\ &= (1-r)\theta_t + \left(\frac{1}{1-\alpha} - r \right) u_t^x - \alpha u_t^n. \end{aligned}$$

The first equality in (38) shows that output choices are strategic substitutes: when other firms raise their output y_t , each individual firm's output choice responds negatively (since $\sigma > 1$). Indeed, the increase in the real wage when output is high dominates the increase in demand in (33) for all $\sigma > 1$. By contrast, when $\sigma \rightarrow 1$ the demand and real wage effects precisely offset each other, and firms act independently of one another. The second equality in (38) expresses the same relationship in equilibrium, in terms of the latent factor θ_t and other primitive shocks. We conclude that, due to strategic substitutability, the aggregate supply component is *countercyclical*, since it places a negative weight $(1-r) < 0$ on the latent factor. This type of strategic substitutability (or “general equilibrium offset”) arises commonly in flexible-price business cycle models, especially those that generate realistic amounts of volatility in hours worked (Hansen 1985, Rogerson 1988), because increases in other firms' outputs tend to drive up production costs. For example, we note that strategic substitutes are also a key feature of the standard, perfectly competitive real business cycle (RBC) model. Indeed, equations (36) to (38) collapse to the output choice of firms in a standard RBC model (with logarithmic utility from consumption, linear disutility from labor, and no capital accumulation) in the limit as $\sigma \rightarrow \infty$, which clearly corresponds to strategic substitutes.

In online Appendix I, we consider a model that nests both our example and the model in Angeletos and La'O (2010) and Angeletos, Iovino, and La'O (2016). In this extension, among other additional parameters, households have a flexible coefficient ψ of relative risk aversion (our model fixes $\psi = 1$). We show that output choices are strategic substitutes if and only if $\sigma\psi > 1$. Common values in macroeconomics for σ and ψ are $\sigma \geq 4$ and $\psi \geq 1$ (e.g., Galí 2008, chap. 2). Hence, while qualitative explorations of models of strategic complementarity have yielded important insights (e.g., Angeletos, Iovino, and La'O 2016), we view the case in which output choices are strategic substitutes as a quantitatively relevant one for this class of models.

The properties above, along with our results in Proposition 2 and 3, suggest that firms' expectations about future output will match the characteristics of the survey data when firms pay imperfect, asymmetric attention to the first component x_{1t} . For

example, consider the case in which all firms except firm i pay full attention to both components, while firm i pays full attention to x_{1t} but none to x_{2t} . Then, it immediately follows that the slope coefficient in a regression of firm i 's forecast errors on recent output (that is, similar to (1)) becomes

$$(39) \quad \gamma_i = \text{cov}(y_{t+1} - E_{it}y_{t+1}, y_t) \text{var}[y_t]^{-1} = -\rho \frac{\alpha}{1-\alpha} \frac{\text{var}_t[\theta_t]}{\text{var}[y_t]} < 0,$$

so that firm i appears to extrapolate current output into the future.³²

Equilibrium with Limited Attention.—We now return to the full model with limited attention. We start by describing firms' optimal output choices under limited attention, and their corresponding expected profits.

PROPOSITION 5: *An individual firm's output choice under limited attention satisfies $y_{it} = E_{it}[y_{it}^*] = E_{it}[x_{i1t} + x_{2t}]$, and the associated expected, discounted profits are $v_{it}^* \simeq -(1/2)\text{var}[y_{it}^*|\Omega_{it}]$.*

The characterization in Proposition 5 follows from Proposition 4. It further implies that, as a function of attention choices, the unconditional expectation of firms' profit is constant across time. This allows us to state an individual firm's attention choice problem as the following static problem: in the initial period $t = 0$, the firm chooses attention coefficients m_1 and m_2 to maximize

$$(40) \quad \max_{\{m_1, m_2\} \in [0,1]^2} -\frac{1}{2} \text{var}[y_{it}^*|\Omega_{it}] - K(m),$$

while anticipating that its optimal output choice in the subsequent periods will be

$$(41) \quad y_{it} = E[y_{it}^*|z_{i1}^t, z_{i2}^t] = E_{it}[x_{i1t} + x_{2t}],$$

where x_{2t} depends upon $y_t = \int_0^1 y_{it} di$. Notice that the problem in (40) and (41) is an application of the problem we studied in Section III. There are $N = 2$ components of output, which determine the firm's ideal action y_{it}^* . The weight on each component x_{jt} is one ($w_j = 1$). A small modification is that, due to firm-specific shocks, the ideal output y_{it}^* is now firm specific.³³

Numerical Solution Method.—Unlike the full-attention version of the model, the equilibrium dynamics of output can no longer be derived analytically when firms

³²This follows from

$$\begin{aligned} \gamma_i &= \text{cov}(y_{t+1} - E_{it}y_{t+1}, y_t) \text{var}[y_t]^{-1} = \text{cov}\left[y_{t+1} - E_{it}y_{t+1}, x_{1t} \pm \frac{1}{r}\left(\frac{1}{1-\alpha} - r\right)x_{1t}\right] \text{var}[y_t]^{-1} \\ &= \rho \text{cov}\left[\theta_t - E_{it}\theta_t, (1-r)\theta_t - \left(\frac{1}{1-\alpha} - r\right)\theta_t\right] \text{var}[y_t]^{-1} \\ &= -\rho \frac{\alpha}{1-\alpha} \text{var}_t[\theta_t] \text{var}[y_t]^{-1} < 0. \end{aligned}$$

³³Nevertheless, from a firm's perspective, firm-specific shocks are equivalent to an increase in the volatility of component-specific disturbances. Hence, the same conditions as in Section III apply here.

pay limited attention. Instead, we solve the model numerically, looking for linear equilibria in which the law of motion for the components and the latent factor take the form of an infinite dimensional vector,

$$(42) \quad \mathbf{x}_t = A\mathbf{x}_{t-1} + Bu_t, \quad u_t = [u_t^\theta \ u_t^x \ u_t^n]',$$

where $\mathbf{x}_t = [\bar{x}'_{t-1} \ \bar{x}'_{t-2} \ \cdots]'$ with $\bar{x}_t = [x_{1t} \ x_{2t} \ \theta_t]'$ and $x_{1t} = \int_0^1 x_{it} di$, and where A and B are matrices of undetermined coefficients whose rows conform with (28) and (33).

To solve for the rational expectations equilibrium, we further conjecture that

$$(43) \quad y_t = \bar{E}_t[x_{1t} + x_{2t}] = [1 \ 1 \ 0] \bar{E}_t[\mathbf{x}_t] = [1 \ 1 \ 0] \Xi \mathbf{x}_t,$$

where Ξ is another matrix of undetermined coefficients.

Solving the model requires finding values for the matrices A , B , and Ξ , as well as firms' attention choices $m = [m_1 \ m_2]$, which are consistent with firm optimality, Bayesian updating of expectations, and market clearing. We do so by first truncating the infinite-dimensional vector \mathbf{x}_t . In accordance, with Hellwig and Venkateswaran (2009) and Lorenzoni (2009), we truncate it at \bar{x}_{t-T} where $T = 50$, but our numerical results are already stable from around $T = 10$. We then iterate on the following two steps until convergence.

First, we hold attention choices m fixed and derive new matrices A , B , and Ξ implied by Bayesian updating and firm optimality. Specifically, we solve firms' signal extraction problem using the Kalman filter, which implies a new matrix Ξ , characterizing average expectations about \mathbf{x}_t . This matrix, along with firms' optimality conditions, implies new matrices A and B characterizing the law of motion for \mathbf{x}_t , which in turn implies a new matrix Ξ . We iterate on these updates until the coefficients in A , B , and Ξ converge in the sense of absolute difference.

Second, we hold coefficients in A , B , and Ξ fixed and derive new values m for firms' optimal attention choices. We derive an expression for firms' profits in (40) as a function of attention choices, which closely resembles the expression in Lemma 2. We then find new optimal choices m by solving the problem in (40). We halt the iteration between these two steps when attention choices m have converged in the sense of absolute difference. Online Appendix J contains further details about the solution method and its implementation.

C. A Quantitative Exploration

We now explore the quantitative implications of the model. We address two questions: First, can the model match the magnitude of extrapolation and underreaction from the survey data? Second, if so, what are the implications for the dynamics of output? To tackle these questions, we parameterize the model and compare estimates of (1) and (2) to those from the data.

Calibration.—We set the labor share $\alpha = 2/3$ and elasticity of substitution $\sigma = 10$. The persistence of the latent factor θ_t is set to $\rho = 0.90$ and the standard deviation of the shock to $\sigma_\theta = 0.75$. The standard deviation of the transitory

component of productivity is set to $\sigma_x = 1.25$, while the standard deviation of the labor supply shock is set to $\sigma_n = 0.1$. These values are all within the range used in standard dynamic stochastic general equilibrium models with monopolistic competition. Our baseline calibration eliminates firm-specific productivity shocks by setting $\sigma_a = 0$, to cleanly illustrate the effect of attention choices without exogenous noise in firms' information. We later explore the robustness of our results towards this assumption.

For the attention cost function, we use the functional form $K(q) = \mu \sum_j q_j^{-2}$; that is, a marginal cost μ multiplied by the sum of signal precisions $1/q_j^2$ across the components of output (Veldkamp 2011).³⁴ The free parameter is the marginal cost μ , which determines the overall imperfection in firms' information. For example, if $\mu = 0$, then we obtain the full information benchmark, because firms can obtain infinitely precise signals at no cost.

As Coibion and Gorodnichenko (2015) points out, information frictions relate directly to the observable coefficient δ in (2) that measures underreactions in average revisions. Hence, we calibrate μ to match estimated underreactions. Concretely, we solve the model repeatedly, varying μ , until the estimate of $\hat{\delta}$ obtained from the model's output matches the empirical estimate obtained from one-quarter-ahead forecasts in the SPF. This approach yields $\mu = 2.00$. This calibration implicitly assumes that forecasts reported by respondents in the SPF are similar to the expectations of firms in our model. Clearly, survey respondents may instead be motivated by career concerns, a desire to attract publicity, or other biased incentives (e.g., Ehrbeck and Waldmann 1996, Lamont 2002, Ottaviani and Sørensen 2006). The related empirical evidence is mixed.³⁵ Following the literature, we view the estimates from professional forecasters as providing a useful lower-bound on deviations from full-information rationality.³⁶

Components of Output and Attention Choices.—Recall from Proposition 2 and 3 that (i) asymmetric attention to procyclical variables can rationalize apparent extrapolation and underreactions, and that (ii) these patterns are consistent with optimal attention choices if procyclical variables are either more volatile or more important for agents' decision-making. Figure 6 and Table 2 illustrate these mechanisms in general equilibrium.

Figure 6 shows that, as in the full information case, the productivity component is procyclical, while the aggregate supply component is countercyclical in equilibrium. Output as a whole is procyclical. The first two columns in Table 2 show the significance of the productivity and the aggregate supply component in firms' decision problem. While both components have a utility weight of one in firms' ideal output choice (Proposition 4), the productivity component is much more

³⁴In equilibrium, there is a one-to-one mapping between the precision parameters q_j and the attention parameters m_j . Similar conclusions as those presented in Table 2 arise with an entropy-based cost function.

³⁵For example, Lamont (2002) finds evidence for strategic forecasts in the *nonanonymized Business Week Survey*, but Stark (1997) argues that the same hypothesis is rejected in the *anonymized* SPF. Ehrbeck and Waldmann (1996) rejects a model of strategically biased forecasts in Treasury bill forecasts from the Blue Chip survey.

³⁶See, for example, Lorenzoni (2009), Nimark (2014), and Angeletos and Huo (2021). We note that the SPF includes forecasts from large industrial firms, in addition to those from financial and government institutions, and forecasting agencies. The biannual Livingston Survey estimates reported in Section I, which resemble those from the SPF, include a broader range of nonfinancial firms.

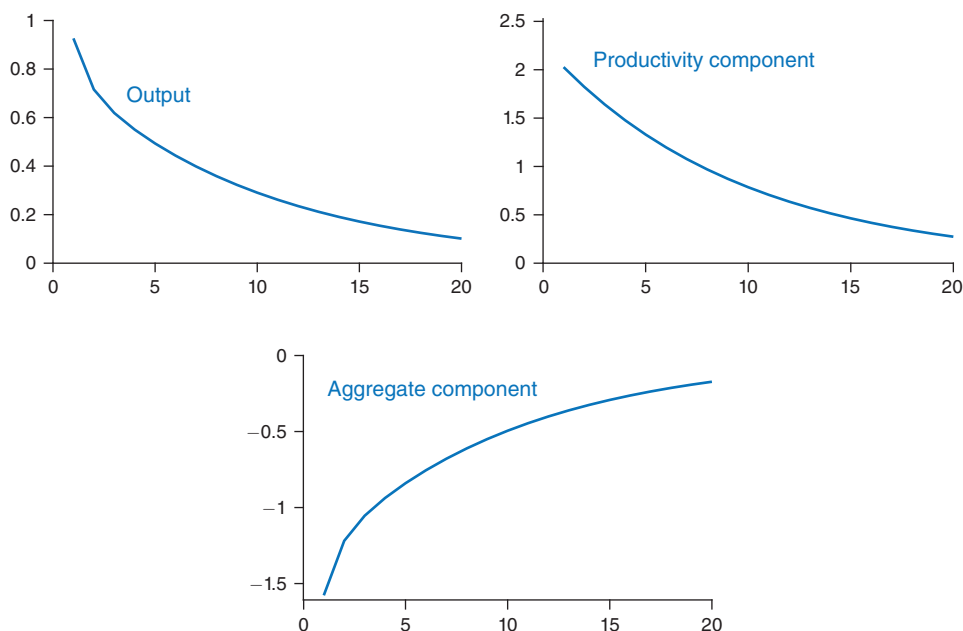


FIGURE 6. CYCLICALITY OF STRUCTURAL COMPONENTS AND OUTPUT

Notes: The figure depicts the impulse response function to a unit standard deviation shock to θ_t on the vertical axis. Time is measured in quarters on the horizontal axis.

TABLE 2—ATTENTION CHOICES IN EQUILIBRIUM

Component	Variance	Weight	q	m
Productivity component (x_{1t})	2.44	1.00	1.29	0.87
Economy-wide component (x_{2t})	0.74	1.00	1.66	0.14

Note: Variances have been scaled by the variance of output.

volatile for baseline parameters.³⁷ The third and fourth columns in Table 2 show firms' optimal attention choices (m_j), or equivalently noise choices (q_j), for both output components. As expected, attention gravitates towards the productivity component x_{1t} because of its larger volatility. In particular, firms optimally choose to pay around four times more attention to x_{1t} . This is consistent with the conclusions from Lucas (1977) and Maćkowiak and Wiederholt (2009) (also cited in the introduction) that for most firms there is little reason to pay particularly close attention to aggregate conditions. Coibion, Gorodnichenko, and Kumar (2018b) provides evidence in favor of this supposition. We now explore the implications of these asymmetries for firms' expectations in equilibrium.

³⁷ Notice that, because firms have imperfect information about both components, the variance of each component in Table 2 can exceed that of output itself (which is the expectation of the sum).

TABLE 3—OVER- AND UNDERREACTIONS

	Model estimates		SPF estimates	
	Forecast error	Forecast error	Forecast error	Forecast error
<i>Panel A. One-quarter-ahead horizon</i>				
Current realization	−0.09 (—)		−0.05 (0.06)	
Average revision		0.38 (—)		0.38 (0.16)
Sample RRMSE	(—)	0.93 (—)	1970IV:2019IV	1970IV:2019IV
<i>Panel B. Four-quarter-ahead horizon</i>				
Current realization	−0.06 (—)		−0.12 (0.05)	
Average revision		0.30 (—)		0.68 (0.19)
Sample RRMSE	(—)	0.92 (—)	1970IV:2019IV	1970IV:2019IV

Notes: Double-clustered robust standard errors in parentheses. The top/bottom 1 percent of forecast errors and revisions has been trimmed pre-estimation. *RRMSE* denotes the root mean-squared error of individual forecasts relative to an estimated AR(1).

Over- and Underreactions.—The first two columns in Table 3 panel A show the results of estimating the extrapolation regression (1) and the underreactions regression (2) on firms' simulated expectations of one-quarter-ahead output in equilibrium. The third and fourth columns compare these estimates to the magnitude of those obtained in the survey data at the one-quarter horizon (Table C.2 in the online Appendix). The underreaction coefficient δ at the one-quarter frequency was a targeted moment. Due to firms' asymmetric attention to the procyclical component of output, the coefficient γ on current output in (1) is negative, generating apparent overreactions in expectations that are close to those in the data. As a result, firms' expectations are simultaneously consistent with extrapolation and underreactions.

Table 3 panel B shows the implied estimates at the four-quarter horizon, which mirror the horizon in Table 1. The model does not match the increase in the magnitude of these coefficients. One difference that may drive this quantitative discrepancy is that the model-implied estimates in Table 3 use forecast errors about the level of output, since our model assumes that output follows a stationary process, while our empirical estimates in Table 1 are based on forecasts about the growth rates of output.³⁸ Despite this simplification, the model estimates are still of a commensurable size to the empirical estimates at the four-quarter horizon, neither of which were targeted moments in the calibration.³⁹

³⁸We note that an alternative calibration of the model that increases the standard deviation of the persistent component to $\sigma_\theta = 1.00$ with $\sigma = 6$ accurately matches the size of the estimates in Table C.2 in the online Appendix using firms' expectations of the growth rate of output instead of the level. The model-implied estimate of μ is 1.30, and the estimates at the four-quarter horizon still decline in magnitude. None of the main conclusions from Section IVC and IVD change under this alternative calibration.

³⁹An alternative approach is to calibrate the model by targeting the four-quarter δ estimate in Table 1. In this case, we arrive at estimates for γ that are close to their empirical counterparts. The implied one-quarter-ahead estimates, however, suggest slightly more extrapolation than what we see in the data.

The last row in Table 3 shows that firms in the simulated model make better forecasts (in a root-mean square error sense) than they would achieve using a simple time-series model. This is consistent with our empirical results in Section III.

D. Further Implications of Asymmetric Attention

We leverage our calibrated model to illustrate two wider implications of asymmetric attention. First, we show that asymmetric attention causes the equilibrium dynamics of output to be more persistent and more volatile. Second, we show that our model is also consistent with increased responsiveness to new information, and increased extrapolation, after the onset of the Great Moderation (as we also document empirically in the online Appendix).

Asymmetric Attention and Output Dynamics.—We compare the dynamics of output in our model with those that arise in an equivalent model where attention is limited but *symmetric*. In this symmetric case, firms observe only one noisy signal of their optimal output:

$$(44) \quad s_{it} = y_t^* + q_* \epsilon_{it} = x_{1t} + x_{2t} + q_* \epsilon_{it}, \quad \epsilon_{it} \sim \mathcal{N}(0, 1),$$

where the noise parameter q_* is calibrated to match firms' uncertainty about their optimal output choice in the asymmetric attention version of the model.

Figure 7 summarizes the results. The left panel shows that the model with asymmetric attention results in more persistence in output (larger autocorrelation). This is intuitive: when firms focus their attention on the procyclical, productivity component, their beliefs and actions become more persistent, because this component directly tracks the dynamics of the latent factor. This increase in persistence occurs even though all input choices happen within the period. An additional, predetermined factor of production, such as capital, would amplify these effects by allowing firms' extrapolative expectations to directly affect future output.

Relatedly, the right panel in Figure 7 shows that output responses are also more correlated with the latent factor itself when there is asymmetric attention. The bottom panel, in turn, shows that asymmetric attention also causes the unconditional variance of output to increase. For the same overall information friction (as measured by δ in (2)), the asymmetry of attention increases the volatility of output, and pushes it closer to its full information value.

Finally, in line with our previous results, we note that the model with symmetric attention produces a positive estimate of γ ($\gamma = 0.15$).

Asymmetric Attention and the Great Moderation.—One manifestation of the Great Moderation was a reduction in the size of aggregate versus firm-specific shocks. As discussed in, for example, Arias, Hansen, and Ohanian (2007) and Galí and Gambetti (2009), the standard deviation of aggregate productivity shocks declined by around 40–50 percent after 1985, while the volatility of firm-specific shocks appears mostly unchanged (Comin and Philippon 2005). We explore the implications of a similar structural shift in our model.




FIGURE 7. ASYMMETRIC ATTENTION AND OUTPUT DYNAMICS

Notes: The left panel shows the autocorrelation of output on the vertical axis, with the lags of output up to four quarters on the horizontal axis. The right panel shows the correlation of output with total factor productivity $a_t = \theta_t + u_t^x$ once more up to a four-quarter lag. We depict these both for the calibrated asymmetric attention model, the symmetric attention model, as well as the full information case. The bottom panel illustrates the variance of output in the asymmetric and symmetric cases relative to the full information benchmark.

Following Arias, Hansen, and Ohanian (2007), we assume that all of the decrease in the volatility of aggregate productivity is due to a decrease in the common, persistent component σ_θ . To model the economy before the Great Moderation, we use our baseline calibration above, but re-introduce firm-specific productivity shocks $\sigma_a > 0$. This parameter is calibrated to match the level of information frictions before the Great Moderation, which we estimate by running regression (2) for one-quarter-ahead forecasts on a sample until 1985:I. To model the economy after the Great Moderation, we then reduce the volatility of σ_θ by 45 percent.

Table 4 shows the resulting estimates of (1) and (2) on model-generated data before and after the Great Moderation. As in the equivalent regressions on the actual survey data, underreactions become weaker while extrapolation becomes somewhat stronger. This is because the decrease in the volatility of common shocks causes firms to choose more asymmetric attention. Indeed, compared to the pre-Great Moderation values, our solution shows that post-Great Moderation firms pay 2 percent more attention to the procyclical component (as measured by q_1), and 27 percent less attention to the countercyclical component.

The results in this subsection have highlighted two implications of asymmetric attention. First, asymmetric attention not only affects the properties of expectations, but also heightens the persistence and volatility of output fluctuations in general equilibrium. Second, an exploration of the Great Moderation provides some

TABLE 4—MODEL ESTIMATES PRE/POST-GREAT MODERATION

	Pre-Great Moderation		Post-Great Moderation	
	Forecast error	Forecast error	Forecast error	Forecast error
Current realization	−0.09 (—)	—	−0.13 (—)	
Average revision		0.55 (—)		0.47 (—)

Notes: Columns 1 and 3 report estimates using one-year-ahead forecast, while columns 2 and 4 employ one-quarter-ahead forecasts. Column 2 is calibrated. The equilibrium noise in signals about the components is pre-Great Moderation $q_1 = 1.43$ and $q_2 = 2.05$, and post-Great Moderation $q_1 = 1.40$ and $q_2 = 2.61$.

validation of our example framework. A simple model based on asymmetric attention to a procyclical, local component of output can qualitatively match the stylized feature that extrapolation strengthened while underreactions subsided at a time when aggregate productivity became less volatile.

V. Conclusion

In this paper, we have contributed to a research agenda that seeks to find a data-consistent model of expectation formation. The framework we have considered relies on minimal frictions relative to the classical benchmark. The only primitive deviation from full information and rational expectations is limited attention. Previous work by Woodford (2001), Sims (2003), Angeletos and Huo (2021), and others, have demonstrated that limited attention offers an explanation for the myopia and anchoring to past outcomes commonly documented in macroeconomics. Our results show that extrapolation, and more generally overreactions to public information, can also be explained by this framework.

We have documented that households', firms', and professional forecasters' expectations simultaneously *overreact* to recent outcomes of the forecasted variable but *underreact* to new information on average. These facts are inconsistent with standard behavioral models of extrapolation, as well as with models that combine the overconfidence inherent to extrapolation with noisy information. To resolve this friction, we have proposed a simple, rational model of limited attention in which people internalize that a forecasted variable is comprised of several components. We characterized the conditions under which this model is consistent with the data. In doing so, we have developed a rational theory of extrapolation that is also consistent with observed underreactions. This theory is based on individuals' *asymmetric attention* to procyclical variables. Through the lens of this model, the overreactions to recent outcomes documented in survey data can be viewed as underreactions to countercyclical components.

To illustrate our results, we embedded our analysis in a workhorse macroeconomic model. For reasonable parameters, we showed that firms' expectations exhibit extrapolation and underreactions, similar to their empirical counterparts. This application also allowed us to study the implications of asymmetric attention for the dynamics of output, and to validate the model further by studying its implications for structural changes around the Great Moderation.

Beyond the analysis in this paper, our results suggest that models of limited, asymmetric attention can account for flexible patterns of predictability in people's forecast errors. We see important scope for extending our results to account for the more general under- and overreactions to public information documented in the literature.⁴⁰ Another avenue for future research is to combine models of optimal information choice with insights from behavioral economics, such as those discussed recently by Bordalo, Gennaioli, and Shleifer (2018). The latter approach would allow for an empirical estimate of the relative contribution of each component to the predictability of forecast errors. Overall, we view the research in this paper as a useful step towards a unified, data-consistent model of expectations based on a minimal set of frictions.

APPENDIX A. PROOFS AND DERIVATIONS

A. Alternative Models

PROOF OF PROPOSITION 1:

The proof proceeds in three steps. We first derive the moving average (MA) form for the nowcast $f_{it}y_t$. We then use this result to derive slope coefficients in (1) and (2).

Step (i): MA form of nowcast. Solving (5) backwards for $k = 0$, we find that

$$(A1) \quad f_{it}y_t = g_0z_{it} + \lambda\rho(1 - g_0)f_{it-1}y_{t-1} = g_0 \sum_{h=0}^{\infty} \lambda^h \rho^h (1 - g_0)^h z_{it-h},$$

where we have also used that $f_{it-1}z_{it} = f_{it-1}y_t = \rho f_{it-1}y_{t-1}$.

Step (ii): Slope coefficient γ in (1). The overreaction coefficient γ equals

$$\begin{aligned} \gamma &= \text{cov}[y_{t+k} - f_{it}y_{t+k}, y_t] \text{var}[y_t]^{-1} \\ &= \rho^k \text{cov}[y_t - f_{it}y_t, y_t] \text{var}[y_t]^{-1} = \rho^k (1 - \text{cov}[f_{it}y_t, y_t] \text{var}[y_t]^{-1}), \end{aligned}$$

because $f_{it}y_{t+k} = \rho^k f_{it}y_t$, and where (A1) shows that

$$\text{cov}[f_{it}y_t, y_t] = g_0 \sum_{h=0}^{\infty} \lambda^h \rho^h (1 - g_0)^h \rho^h \text{var}[y_t] = g_0 \frac{1}{1 - \lambda \rho^2 (1 - g_0)} \text{var}[y_t].$$

Hence,

$$\gamma = \rho^k \left(1 - \frac{g_0}{1 - \lambda \rho^2 (1 - g_0)} \right) = \rho^k (1 - g_0) \frac{1 - \lambda \rho^2}{1 - \lambda \rho^2 (1 - g_0)}.$$

⁴⁰Consider, for example, our baseline model from Section II, and suppose that instead of regression (1) we regress forecast errors onto one component x_{jt} of output. The slope coefficient from this regression would be proportional to $a_j(1 - m_j)$, which could be either positive (representing an underreaction to x_{jt}) or negative (representing an overreaction), depending on the cyclical nature of x_{jt} (the sign of a_j). In principle, we therefore conjecture that the conditions in Proposition 2 could be extended and used to account for the much broader patterns of predictability documented, for example, by Pesaran and Weale (2006) and Fuhrer (2017).

We conclude the sign of γ depends only the sign of $\rho^k(1 - g_0) = \rho^k - g_k$, since the responsiveness coefficient g_k satisfies $g_k = g_0\rho^k$ from (5) and $f_{it}y_{t+k} = \rho^k f_{it}y_t$.

Step (iii): Slope coefficient δ in (2). Averaging (5) across i for $k = 0$, using that $\bar{f}_t y_{t+k} = \rho^k \bar{f}_t y_t$, and rearranging terms as in Coibion and Gorodnichenko (2015) shows that:

$$y_{t+k} - \bar{f}_t y_{t+k} = \frac{1 - g_0}{g_0} \rho^k (\bar{f}_t y_t - \lambda \bar{f}_{t-1} y_t) + u_{t,t+k},$$

where $u_{t,t+k}$ denotes a linear combination of future shocks $(u_{t+s})_{0 \leq s \leq k}$ to output.

Thus, the underreaction coefficient δ equals

$$\begin{aligned} \delta &= \text{cov}[y_{t+k} - f_{it}y_{t+k}, \bar{f}_t y_{t+k} - \bar{f}_{t-1} y_{t+k}] \text{var}[\chi_t]^{-1} \\ &= \text{cov}[y_{t+k} - \bar{f}_t y_{t+k}, \bar{f}_t y_{t+k} - \bar{f}_{t-1} y_{t+k}] \text{var}[\chi_t]^{-1} \\ &= \rho^k \frac{1 - g_0}{g_0} \rho^k \text{cov}[\bar{f}_t y_t - \lambda \bar{f}_{t-1} y_t, \bar{f}_t y_t - \bar{f}_{t-1} y_t] \text{var}[\chi_t]^{-1}, \end{aligned}$$

where $\chi_t \equiv \bar{f}_t y_{t+k} - \bar{f}_{t-1} y_{t+k}$, and the second equality follows from the linearity of the covariance operator, and because the signals in (4) have the same steady-state distribution for all i . We have used that $\bar{f}_t y_{t+k} = \rho^k \bar{f}_t y_t$ for the third equality. Finally, because g_k satisfies $g_k = g_0 \rho^k$, all that remains to show is that $\text{cov}[\bar{f}_t y_t - \lambda \bar{f}_{t-1} y_t, \bar{f}_t y_t - \bar{f}_{t-1} y_t] > 0$.

Multiplying out terms, and using the stationarity of forecasts, we find that

$$\begin{aligned} &\text{cov}[\bar{f}_t y_t - \lambda \bar{f}_{t-1} y_t, \bar{f}_t y_t - \bar{f}_{t-1} y_t] \\ &= (1 + \lambda \rho^2) \text{var}[\bar{f}_t y_t] - \rho(1 + \lambda) \text{cov}[\bar{f}_t y_t, \bar{f}_{t-1} y_{t-1}] \\ &\geq (1 - \rho)(1 - \lambda \rho) \text{var}[\bar{f}_t y_t] > 0, \end{aligned}$$

since $\bar{f}_{t-1} y_t = \rho \bar{f}_{t-1} y_{t-1}$ and $\text{var}[\bar{f}_t y_t] \geq \text{cov}[\bar{f}_t y_t, \bar{f}_{t-1} y_{t-1}] > 0$.⁴¹ We conclude the sign of δ depends only the sign of the sufficient statistic $\rho^k - g_k$. This completes the proof. ■

COROLLARY 1: *Consider the diagnostic expectations model: $f_{it}y_{t+k} = E_{it-1}y_{t+k} + g_k(z_{it} - E_{it-1}y_t)$. Then, the coefficients γ in (1) and δ in (2) both have the same sign as $\rho^k - g_k$.*

⁴¹ It follows from (5) that

$$\bar{f}_t y_t = \rho[1 + \lambda(1 - g_0)]\bar{f}_{t-1} y_{t-1} - \lambda \rho^2(1 - g_0)\bar{f}_{t-2} y_{t-2} + g_0 u_t.$$

Thus,

$$\text{cov}(\bar{f}_t y_t, \bar{f}_{t-1} y_{t-1}) = \rho \frac{1 + \lambda(1 - g_0)}{1 + \lambda \rho(\rho - g_0 \rho)} \text{var}[\bar{f}_t y_t] > 0.$$

PROOF:

The proof follows from Proposition 1. To see this implication, first notice that the diagnostic nowcast error at time t equals

$$\begin{aligned} y_t - f_{it}y_t &= (1 - g_0)(y_t - E_{it-1}y_t) - g_0\epsilon_{it} \\ &= (1 - g_0)\left(y_t - \frac{1}{1 - g_0^*}E_{it}y_t + \frac{g_0^*}{1 - g_0^*}z_{it}\right) - g_0\epsilon_{it} \\ &= (1 - g_0)(1 - g_0^*)^{-1}(y_t - E_{it}y_t) + [g_0^*(1 - g_0)(1 - g_0^*)^{-1} - g_0]\epsilon_{it}, \end{aligned}$$

where the second equality exploits (5) in the rational case, and we let $g_0^* \in (0, 1)$ denote the noisy rational expectation gain on z_{it} . It now follows from $y_{t+k} - f_{it}y_{t+k} = \rho^k(y_t - f_{it}y_t) + u_{t,t+k}$, where $u_{t,t+k}$ denotes a linear combination of future shocks $(u_{t+s})_{0 \leq s \leq k}$ to output, that

$$y_{t+k} - f_{it}y_{t+k} = (1 - g_0)(1 - g_0^*)^{-1}(y_{t+k} - E_{it}y_{t+k}) + \text{t.u.w.},$$

where t.u.w. denotes *terms uncorrelated with* y_t or $\bar{f}_t y_{t+k} - \bar{f}_{t-1} y_{t+k}$, and we have used (3) and $E_{it}y_{t+k} = \rho^k E_{it}y_t$. We conclude $\gamma = (1 - g_0)(1 - g_0^*)^{-1}\gamma_{NRE}$ and $\delta = (1 - g_0)(1 - g_0^*)^{-1}\delta_{NRE}$, where γ_{NRE} and δ_{NRE} denote the over- and underreaction coefficients, respectively, in the noisy rational expectation case. Proposition 1 implies $\gamma_{NRE} > 0$ and $\delta_{NRE} > 0$. Thus, the sign of γ and δ depend only $1 - g_0 = (\rho^k - g_k)\rho^{-k}$, which depends only on $\rho^k - g_k$. ■

B. Asymmetric Attention

PROOF OF LEMMA 1:

The proof follows directly from the derivation of the Kalman gain g_j .

At date t , agent i 's signal z_{ijt} is informationally equivalent to the signal

$$\hat{z}_{ijt} \equiv \frac{z_{ijt}}{a_j} = \theta_t + \frac{1}{a_j}(b_j u_{jt} + q_j \epsilon_{ijt}) \equiv \theta_t + \xi_{ijt},$$

which has precision $\tau_j \equiv \text{var}[\hat{z}_{ijt}|\theta_t]^{-1}$ equal to

$$\tau_j = \frac{a_j^2}{b_j^2 + q_j^2} = \frac{a_j^2}{b_j^2} m_j.$$

The standard formula for Gaussian updating now implies that

$$(A2) \quad E_{it}[\theta_t] = E_{it-1}[\theta_t] + \sum_j \left(\frac{\tau_j}{\bar{\tau} + \sum_k \tau_k} \right) (\hat{z}_{ijt} - E_{it-1}[\hat{z}_{ijt}]),$$

where $\bar{\tau} \equiv \text{var}[\theta_t | \Omega_{it-1}]^{-1}$, while the posterior precision satisfies $\text{var}[\theta_t | \Omega_{it}]^{-1} = \bar{\tau} + \sum_k \tau_k$.

Combining terms, and inserting the definition of \hat{z}_{ijt} into (A2), we obtain that

$$E_{it}[\theta_t] = E_{it-1}[\theta_t] + \sum_j \text{var}[\theta_t | \Omega_{it}] \frac{a_j}{b_j^2} m_j (z_{ijt} - E_{t-1} z_{ijt}).$$

Equating $g_j = \text{var}[\theta_t | \Omega_{it}] (a_j/b_j^2) m_j$ then completes the proof. ■

PROOF OF PROPOSITION 2:

We start with the characterization of the extrapolation coefficient γ in (1). Equation (12) shows that the sign of γ is determined by

$$\begin{aligned} \text{(A3)} \quad \gamma &\propto \sum_j \text{cov}[\theta_t - E_{it} \theta_t, x_{jt}] = \sum_j (a_j \text{cov}[\theta_t - E_{it} \theta_t, \theta_t] + b_j \text{cov}[\theta_t - E_{it} \theta_t, u_{jt}]) \\ &= \sum_j (a_j \text{var}[\theta_t | \Omega_{it}] - b_j \text{cov}[E_{it} \theta_t, u_{jt}]), \end{aligned}$$

since $\text{cov}(\theta_t, u_{jt}) = 0$ and $\text{cov}[\theta_t - E_{it} \theta_t, \theta_t] = E[(\theta_t - E_{it} \theta_t)^2] = \text{var}[\theta_t | \Omega_{it}]$.

Lemma 1 now implies that

$$\text{cov}[E_{it} \theta_t, u_{jt}] = \text{cov}[g_j z_{ijt}, u_{jt}] = g_j b_j = \text{var}[\theta_t | \Omega_{it}] \frac{a_j}{b_j} m_j.$$

Substituting this expression into (A3), we conclude that

$$\gamma \propto \sum_j \text{cov}[\theta_t - E_{it} \theta_t, x_{jt}] = \text{var}[\theta_t | \Omega_{it}] \sum_j a_j (1 - m_j).$$

This completes the first step of the proposition.

Turning to the characterization of the underreaction coefficient δ in (2), we start by solving the Kalman filter in (11) backwards to obtain

$$\text{(A5)} \quad E_{it}[\theta_t] = \sum_{h=0}^{\infty} \lambda^h \hat{z}_{it-h},$$

where we define the precision-weighted signal $\hat{z}_{it} \equiv \sum_j g_j z_{ijt}$, and let $\lambda \equiv (1 - \sum_j g_j a_j) \rho$. The average precision-weighted signal is $\int_0^1 \hat{z}_{it} di = \hat{z}_{it} - \hat{\epsilon}_{it}$ for all $i \in [0, 1]$ with $\hat{\epsilon}_{it} \equiv \sum_j g_j q_j \epsilon_{ijt}$.

We thus find that the average forecast revision equals

$$\begin{aligned} \bar{E}_t \theta_t - \bar{E}_{t-1} \theta_t &= \bar{E}_t \theta_t - \rho \bar{E}_{t-1} \theta_{t-1} \\ &= \sum_{h=0}^{\infty} \lambda^h (\hat{z}_{it-h} - \hat{\epsilon}_{it-h}) - \rho \sum_{h=1}^{\infty} \lambda^{h-1} (\hat{z}_{it-h} - \hat{\epsilon}_{it-h}). \end{aligned}$$

By the projection theorem, agent i 's forecast error $\theta_t - E_{it}\theta_t$ is uncorrelated with \hat{z}_{it-h} for all $h \geq 0$. Thus, the characterization of δ in (13) yields

$$\begin{aligned}\delta &\propto \text{cov}\left[\theta_t - E_{it}\theta_t, \bar{E}_t\theta_t - \bar{E}_{t-1}\theta_t\right] \\ &= \text{cov}\left[\theta_t - E_{it}\theta_t, -\sum_{h=0}^{\infty}\lambda^h\hat{\epsilon}_{it-h} + \rho\sum_{h=1}^{\infty}\lambda^{h-1}\hat{\epsilon}_{it-h}\right] \\ &= \text{cov}\left[\sum_{h=0}^{\infty}\lambda^h\hat{z}_{it-h}, \sum_{h=0}^{\infty}\lambda^h\hat{\epsilon}_{it} - \rho\sum_{h=1}^{\infty}\lambda^{h-1}\hat{\epsilon}_{it-h}\right] \\ &= \left[1 + \sum_{h=1}^{\infty}\lambda^h(\lambda^h - \rho\lambda^{h-1})\right]\text{var}[\hat{\epsilon}_{it}] \\ &= \frac{1 - \lambda\rho}{1 - \lambda^2}\text{var}[\hat{\epsilon}_{it}],\end{aligned}$$

where the second and third equality use $\text{cov}[\theta_t, \hat{\epsilon}_{it-h}] = 0$ and $\text{cov}[\hat{z}_{it-\ell}, \hat{\epsilon}_{it-h}] = 1_{\ell=h}\text{var}[\hat{\epsilon}_{it}]$.

Since $\lambda < \rho \leq 1$, we conclude that

$$\delta \propto \text{var}[\hat{\epsilon}_{it}] = \sum_j g_j^2 q_j^2 = \text{var}[\theta_t | \Omega_{it}] \sum_j \frac{a_j^2}{b_j^2} m_j (1 - m_j).$$

This expression is positive whenever $0 < m_j < 1$ for at least one j . ■

REFERENCES

- Andrade, Philippe, and Hervé Le Bihan.** 2013. "Inattentive Professional Forecasters." *Journal of Monetary Economics* 60 (8): 967–82.
- Andreassen, Paul, and Stephen Kraus.** 1988. "Judgmental Prediction by Extrapolation." Unpublished.
- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas.** 2018. "Quantifying Confidence." *Econometrica* 86 (5): 1689–1726.
- Angeletos, George-Marios, and Zhen Huo.** 2021. "Myopia and Anchoring." *American Economic Review* 111 (4): 1166–1200.
- Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry.** 2020. "Imperfect Macroeconomic Expectations: Evidence and Theory." NBER Working Paper 27308.
- Angeletos, George-Marios, Luigi Iovino, and Jennifer La'O.** 2016. "Real Rigidity, Nominal Rigidity, and the Social Value of Information." *American Economic Review* 106 (1): 200–227.
- Angeletos, George-Marios, and Jennifer La'O.** 2010. "Noisy Business Cycles." In *NBER Macroeconomics Annual 2009*, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 319–78. Chicago: University of Chicago Press.
- Angeletos, George-Marios, and Jennifer La'O.** 2013. "Sentiments." *Econometrica* 81 (2): 739–79.
- Arias, Andres, Gary D. Hansen, and Lee E. Ohanian.** 2007. "Why Have Business Cycle Fluctuations Become Less Volatile?" *Economic Theory* 32 (1): 43–58.
- Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer.** 2016. "Extrapolation and Bubbles." NBER Working Paper 21944.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny.** 1998. "A Model of Investor Sentiment." *Journal of Financial Economics* 49 (3): 307–43.
- Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni.** 2013. "News, Noise, and Fluctuations: An Empirical Exploration." *American Economic Review* 103 (7): 3045–70.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer.** 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review* 110 (9): 2748–82.

- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer.** 2018. "Diagnostic Expectations and Credit Cycles." *Journal of Finance* 73 (1): 199–227.
- Broer, Tobias, and Alexandre Kohlhas.** 2019. "Forecaster (Mis-)Behavior." Unpublished.
- Cagan, Philip.** 1956. "The Monetary Dynamics of Hyperinflation." In *Studies in the Quantity Theory of Money*, edited by Milton Friedman. Chicago: University of Chicago Press.
- Chahrour, Ryan, and Robert Ulbricht.** 2018. "Robust Predictions for DSGE Models with Incomplete Information." Unpublished.
- Coibion, Olivier, and Yuriy Gorodnichenko.** 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120 (1): 116–59.
- Coibion, Olivier, and Yuriy Gorodnichenko.** 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8): 2644–78.
- Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar.** 2018. "How Do Firms Form Their Expectations? New Survey Evidence." *American Economic Review* 108 (9): 2671–2713.
- Comin, Diego, and Thomas Philippon.** 2005. "The Rise in Firm-Level Volatility: Causes and Consequences." In *NBER Macroeconomics Annual 2005*, edited by Mark Gertler and Kenneth Rogoff, 167–201. *NBER Macroeconomics Annual*, Vol. 20. Cambridge, MA: MIT Press.
- Cover, Thomas M., and Joy A. Thomas.** 2012. *Elements of Information Theory*. Hoboken, NJ: John Wiley & Sons.
- Croushore, Dean D.** 1993. "Introducing: The Survey of Professional Forecasters." *Business Review: Federal Reserve Bank of Philadelphia* 6: 3–15.
- Croushore, Dean D.** 1997. "The Livingston Survey: Still Useful After All These Years." *Business Review: Federal Reserve Bank of Philadelphia* 2: 1–12.
- Croushore, Dean, and Tom Stark.** 2019. "Fifty Years of the Survey of Professional Forecasters." *Economic Insights* 4 (4): 1–11.
- Curtin, Richard T.** 1982. "Indicators of Consumer Behavior: The University of Michigan Surveys of Consumers." *Public Opinion Quarterly* 46 (3): 340–52.
- Cutler, David M., James M. Poterba, and Lawrence H. Summers.** 1990. "Speculative Dynamics and the Role of Feedback Traders." *American Economic Review* 80 (2): 63–68.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam.** 1998. "Investor Psychology and Security Market Under- and Overreactions." *Journal of Finance* 53 (6): 1839–85.
- de Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann.** 1990. "Positive Feedback Investment Strategies and Destabilizing Rational Speculation." *Journal of Finance* 45 (2): 379–95.
- Ehrbeck, Tilman, and Robert Waldmann.** 1996. "Why Are Professional Forecasters Biased? Agency versus Behavioral Explanations." *Quarterly Journal of Economics* 111 (1): 21–40.
- Eyster, Erik, Matthew Rabin, and Dimitri Vayanos.** 2019. "Financial Markets Where Traders Neglect the Informational Content of Prices." *Journal of Finance* 74 (1): 371–99.
- Fuhrer, Jeffrey C.** 2017. "Intrinsic Expectations Persistence: Evidence from Professional and Household Survey Expectations." Unpublished.
- Fuster, Andreas, Benjamin Hebert, and David Laibson.** 2012. "Natural Expectations, Macroeconomic Dynamics, and Asset Pricing." *NBER Macroeconomics Annual 2011*, edited by Daron Acemoglu and Michael Woodford, 1–48. *NBER Macroeconomics Annual*, Vol. 26. Chicago: University of Chicago Press.
- Gabaix, Xavier.** 2017. "Behavioral Inattention." NBER Working Paper 24096.
- Gali, Jordi.** 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton, NJ: Princeton University Press.
- Gali, Jordi, and Luca Gambetti.** 2009. "On the Sources of the Great Moderation." *American Economic Journal: Macroeconomics* 1 (1): 26–57.
- Garcia, Juan A.** 2003. "An Introduction to the ECB's Survey of Professional Forecasters." European Central Bank Occasional Paper 8.
- Goodwin, R. M.** 1947. "Dynamical Coupling with Especial Reference to Markets Having Production Lags." *Econometrica* 15: 181–204.
- Greenwood, Robin, and Samuel G. Hanson.** 2015. "Waves in Ship Prices and Investment." *Quarterly Journal of Economics* 130 (1): 55–109.
- Greenwood, Robin, and Andrei Shleifer.** 2014. "Expectations of Returns and Expected Returns." *Review of Financial Studies* 27 (3): 714–46.
- Hansen, Gary D.** 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics* 16 (3): 309–27.
- Hellwig, Christian, Sebastian Kohls, and Laura Veldkamp.** 2012. "Information Choice Technologies." *American Economic Review* 102 (3): 35–40.

- Hellwig, Christian, and Venky Venkateswaran. 2009. "Setting the Right Prices for the Wrong Reasons." *Journal of Monetary Economics* 56: S57–77.
- Hirshleifer, David, Sonya S. Lim, and Siew Hong Teoh. 2011. "Limited Investor Attention and Stock Market Misreactions to Accounting Information." *Review of Asset Pricing Studies* 1 (1): 35–73.
- Kahneman, Daniel, and Amos Tversky. 1972. "Subjective Probability: A Judgment of Representativeness." *Cognitive Psychology* 3 (3): 430–54.
- Kohlhas, Alexandre N., and Ansgar Walther. 2021. "Replication Data for: Asymmetric Attention." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E135781V1>.
- Kydland, Finn E., and Edward C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (6): 1345–70.
- Lamont, Owen A. 2002. "Macroeconomic Forecasts and Microeconomic Forecasters." *Journal of Economic Behavior and Organization* 48 (3): 265–80.
- Lorenzoni, Guido. 2009. "A Theory of Demand Shocks." *American Economic Review* 99 (5): 2050–84.
- Lucas, Robert E., Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (2): 103–24.
- Lucas, Robert E., Jr. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63 (3): 326–34.
- Lucas, Robert E., Jr. 1977. "Understanding Business Cycles." *Carnegie-Rochester Conference Series on Public Policy* 5: 7–29.
- Lütkepohl, Helmut. 2007. *New Introduction to Multiple Time Series Analysis*. Berlin: Springer.
- Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2018. "Dynamic Rational Inattention: Analytical Results." *Journal of Economic Theory* 176: 650–92.
- Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99 (3): 769–803.
- Maćkowiak, Bartosz, and Mirko Wiederholt. 2015. "Business Cycle Dynamics under Rational Inattention." *Review of Economic Studies* 82 (4): 1502–32.
- Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.
- Moore, Don A., and Paul J. Healy. 2008. "The Trouble with Overconfidence." *Psychological Review* 115 (2): 502–17.
- Muth, J. F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica* 29: 315–35.
- Nimark, Kristoffer. 2008. "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics* 55 (2): 365–82.
- Nimark, Kristoffer P. 2014. "Man-Bites-Dog Business Cycles." *American Economic Review* 104 (8): 2320–67.
- Ottaviani, Marco, and Peter Norman Sørensen. 2006. "The Strategy of Professional Forecasting." *Journal of Financial Economics* 81 (2): 441–66.
- Pesaran, M. Hashem. 1987. *The Limits to Rational Expectations*. Hoboken, NJ: Blackwell.
- Pesaran, M. Hashem, and Martin Weale. 2006. "Survey Expectations." In *Handbook of Economic Forecasting*, Vol. 1, edited by Graham Elliott, Clive W. J. Granger, and Allan Timmermann, 715–76. Amsterdam: Elsevier.
- Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics* 21 (1): 3–16.
- Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.
- Stark, Tom. 1997. "Macroeconomic Forecasts and Microeconomic Forecasters in the Survey of Professional Forecasters." Federal Reserve Bank of Philadelphia Working Paper 97-10.
- Stark, Tom. 2010. *Realistic Evaluation of Real-Time Forecasts in the Survey of Professional Forecasters*. Federal Reserve Bank of Philadelphia Research Rap, Special Report 1.
- Veldkamp, Laura L. 2011. *Information Choice in Macroeconomics and Finance*. Princeton, NJ: Princeton University Press.
- Vives, Xavier, and Liyan Yang. 2018. "Costly Interpretation of Asset Prices." Unpublished.
- Wiederholt, Mirko. 2010. "Rational Inattention." In *The New Palgrave Dictionary of Economics, Online Edition*, edited by Matias Vernengo, Esteban Perez Caldentey, and Barkley J. Rosser Jr. Berlin: Springer Nature.
- Woodford, Michael. 2001. "Imperfect Common Knowledge and the Effects of Monetary Policy." NBER Working Paper 8673.
- Zarnowitz, Victor. 1985. "Rational Expectations and Macroeconomic Forecasts." *Journal of Business and Economic Statistics* 3 (4): 293–311.