

OVERREACTION IN EXPECTATIONS: EVIDENCE AND THEORY*

Hassan Afrouzi
Columbia

Spencer Y. Kwon
Harvard

Augustin Landier
HEC Paris

Yueran Ma
Chicago Booth

David Thesmar
MIT Sloan

February 7, 2023

Abstract

We investigate biases in expectations across different settings through a large-scale randomized experiment where participants forecast stable stochastic processes. The experiment allows us to control forecasters' information sets as well as the data generating process, so we can cleanly measure biases in beliefs. We report three facts. First, forecasts display significant overreaction to the most recent observation. Second, overreaction is stronger for less persistent processes. Third, overreaction is also stronger for longer forecast horizons. We develop a tractable model of expectations formation with costly processing of past information, which closely fits the empirical facts. We also perform additional experiments to test the mechanism of the model. JEL Codes: C91, D84, G41.

*This paper is a combination of two previous works "Biases in Expectations: Experimental Evidence" and "A Model of Costly Recall." Many thanks to Nick Barberis, Cary Frydman, Nicola Gennaioli, David Hirshleifer, Charlie Nathanson, Martin Schneider, Andrei Shleifer, Dmitry Taubinsky, Emanuel Vespa, and Mike Woodford, as well as conference participants at the AFA Annual Meeting and the Behavioral Approaches to Financial Decision Making Conference and seminar participants at Babson College, Berkeley-Chicago Behavioral Econ Workshop, Bocconi, BYU Marriott, Columbia, Duke, EDHEC, EIEF, the Fed, HEC Paris, Yale SOM, MIT Sloan, NBER, Northwestern Kellogg, NYU Stern, Singapore Management University, Stanford SITE, University of Amsterdam, USC Marshall, and Virtual Finance Workshop. David Norris provided very skillful research assistance on this project. Kwon is grateful for support from the Alfred P. Sloan Foundation Pre-doctoral Fellowship in Behavioral Macroeconomics, awarded through the NBER. Landier acknowledges financial support from the European Research Council under the European Community's Seventh Framework Program (FP7/2007-2013) Grant Agreement no. 312503 SolSys. This study is registered in the AEA RCT Registry and the identifying number is: "AEARCTR-0003173." To preview or set up the experimental interface, please visit <https://github.com/forecast-research/interface>. Emails: hassan.afrouzi@columbia.edu, ykwon@hbs.edu, landier@hec.fr, yueran.ma@chicagobooth.edu, thesmar@mit.edu.

1 Introduction

A growing body of research using survey data shows that survey expectations exhibit significant biases. Across different settings, however, the biases seem to vary. For instance, some studies document substantial overreaction, whereas others find less overreaction or some degree of underreaction.¹ Why do biases in expectations vary across settings? This paper presents new empirical evidence on how expectation biases vary with the persistence of the data generating process (DGP) and the forecast horizon. Based on the empirical findings, we then examine theoretical frameworks of expectation formation.

To cleanly document the biases in expectations, we begin with a large-scale randomized forecasting experiment. Our experimental approach allows us to address three important limitations for studying expectations using survey data. First, we can control forecasters' information sets, which are not observable to the econometrician in survey data. Second, we know and can specify the DGP, whereas the DGP is difficult for the econometrician to pin down in survey data. Finally, we can also control forecasters' payoff functions, whereas forecasters could have considerations other than accuracy in survey data.

In our experiment, participants make forecasts of simple AR(1) processes. They are randomly assigned to a condition with a given AR(1) process, where the persistence ρ is drawn from $\{0, 0.2, 0.4, 0.6, 0.8, 1\}$; the mean is zero, and the conditional volatility is 20. Participants observe 40 past realizations at the beginning and then make forecasts for another 40 rounds. In each round, participants observe a new realization and report one- and two-period-ahead forecasts. In follow-up experiments, we also elicit five- and ten-period-ahead forecasts.

We document three empirical facts. First, even though the process is simple and stable, rational expectations are strongly rejected in our data. In particular, forecasts in the data display significant overreaction to recent observations: the forecasts are systematically too high when the past realization is high, and vice versa. This feature is robust and it does not depend on whether participants know the process is AR(1), which we show using a sample

¹For overreaction in expectations, see [De Bondt and Thaler \(1990\)](#), [Amromin and Sharpe \(2013\)](#), [Greenwood and Shleifer \(2014\)](#), [Gennaioli, Ma and Shleifer \(2016\)](#), [Bordalo, Gennaioli, La Porta and Shleifer \(2019\)](#), [Bordalo, Gennaioli, Ma and Shleifer \(2020c\)](#), [Barrero \(2021\)](#), among others for evidence from forecasts of financial market and macroeconomic outcomes. For underreaction, see [Abarbanell and Bernard \(1992\)](#), [Bouchaud, Krueger, Landier and Thesmar \(2019\)](#), and [Ma, Ropele, Sraer and Thesmar \(2020\)](#) for evidence from forecasts of companies' near-term earnings.

of MIT students who understand AR(1) processes.

Second, we find that forecasts display more overreaction when the process is more transitory. This result echoes the patterns [Bordalo et al. \(2020c\)](#) observe in survey data. In the experiment, however, we can measure the degree of overreaction more precisely. Specifically, we calculate the persistence implied by participants' forecasts (i.e., the regression coefficient of the forecast $F_t x_{t+1}$ on x_t) and compare it with the actual persistence of the process. In our setting, the gap between the implied persistence and the actual persistence provides a clear measure of overreaction. We find that the implied persistence is close to one when the process is a random walk. When the actual persistence decreases, the implied persistence decreases but less than one for one, so the gap between implied and actual persistence is larger when the process is more transitory. For instance, the implied persistence is 0.85 when the actual persistence is 0.6 and 0.45 when the actual process is i.i.d. (i.e., more overreaction when the process persistence is lower).

Third, we find that overreaction is stronger for longer-horizon forecasts. The forecast-implied persistence (per period) is higher for long-horizon forecasts than for short-horizon forecasts, based on forecasts the same participants made for the same AR(1) process.² Accordingly, biases arising from forecasters using a given incorrect value of the persistence parameter (e.g., [Gabaix, 2018](#)) are not sufficient to account for the behavior of forecasts across different horizons. Our finding aligns with the evidence in [Giglio and Kelly \(2018\)](#), who show that affine asset pricing models using a given process persistence cannot simultaneously account for the price movements of short-maturity and long-maturity claims, and the long end implies much higher persistence.

We then examine how well commonly-used expectations models explain the empirical evidence. The older adaptive or extrapolative models generate forecast-implied persistence that does not vary much with the actual persistence, so overreaction in these models is too strong for transitory processes. Some recent models such as sticky and noisy information models ([Mankiw and Reis, 2002](#); [Woodford, 2003](#)) do not generate overreaction, and others such as constant gain learning of the process ([Evans and Honkapohja, 2001](#)) generate implied persistence that aligns too much with the actual persistence, so overreaction is too

²For general forecast horizons, we compute the implied persistence per period by taking the regression coefficient of the forecast $F_t x_{t+h}$ on x_t and raising it to the $1/h$ -th power.

weak for transitory processes.

Taken together, the evidence of overreaction points to models where positive shocks bring to mind higher outcomes that get overweighted in expectations (e.g. [Bordalo, Gennaioli and Shleifer, 2018](#)). Moreover, for transitory shocks to generate overreaction (like what we observe in the data), it is important to allow biases to affect the assessment of the more persistent aspects of the process. Although the standard specification of diagnostic expectations in [Bordalo, Gennaioli and Shleifer \(2018\)](#) does not generate overreaction to transitory shocks since higher future outcomes are not more likely given transitory shocks, [Bordalo et al. \(2019\)](#) feature diagnostic agents who erroneously attribute part of the transitory shock to a more persistent component of the process. Relatedly, [da Silveira and Woodford \(2019\)](#) and [Nagel and Xu \(2022\)](#) model overreaction through biases about the long-run mean.

Building on this insight, we find that a parsimonious way to account for the evidence is to allow recent observations to influence the assessment of the long-run mean of the process (i.e., the most persistent aspect). Based on the psychological foundations of how information comes to mind, we present a tractable model that microfound the degree of overreaction, yields closed-form solutions that are easy to estimate in the data, and generates additional testable predictions. In the model, agents form beliefs about the long-run mean of the process using recent data along with additional information obtained by costly information processing. We draw on two findings from the psychology literature and in particular the field of working memory. First, among available information, agents use a subset that is in a state of “heightened activation” ([Baddeley and Hitch, 1993](#); [Cowan, 1998, 2017a](#)). We refer to the information with heightened activation as “on top of mind.” Second, information with heightened activation consists of recent data, which is automatically incorporated due to the recency effect (i.e., “the enhanced retention of the final item”), and additional task-relevant data processed with a cost ([Hitch, Hu, Allen and Baddeley, 2018](#)). Our model nests the rational benchmark in the frictionless limit where all information is on top of mind. Otherwise, the model delivers partial dependence of the long-run mean assessment on the recent observation. Accordingly, forecasts display overreaction, and overreaction is stronger when the process is less persistent and the forecast horizon is longer.

We take our model to the data by minimizing the mean squared error with respect to all one-period-ahead forecasts in our baseline experiment (we test other models in the litera-

ture in the same way). We calculate the implied persistence in the experimental data and the value generated by the model. The implied persistence based on our model closely matches that in the data for all values of ρ . We then use the model estimated on one-period-ahead forecasts to compute the implied persistence for long-horizon forecasts as non-targeted moments. For forecast horizons of two, five, and ten covered by our experiments, our model closely matches the degree of overreaction observed in the data for each values of ρ .

We also design additional experiments to test the mechanisms of the model. To examine overreaction in the long-run mean belief, we perform experiments where in each round participants enter their long-run mean assessment into a box, by responding to the question “what do you think is the long-run average of this process?” Participants also make long-horizon forecasts about x_{t+10} in each round, so we can compare the consistency between the long-run mean beliefs and the x_{t+10} forecasts.

We find that the long-run mean assessment has a significant positive loading on the most recent observation. We also find that the long-run mean assessment helps explain the x_{t+10} forecast. To examine whether overreaction is influenced by the extent to which the latest information is on top of mind, we perform experiments where we draw participants’ focus away from the most recent observation. In one condition, we require participants to click on the realization ten periods ago before they can make new forecasts in each round; in another condition, we draw a red line at zero in the experimental interface. In both conditions, the degree of overreaction decreases relative to the baseline.

Literature Review. Our work is related to three branches of literature. First, our empirical findings complement recent studies using survey data discussed in the first paragraph, which document strong overreaction in some settings and weaker overreaction or underreaction in others. As mentioned before, while analyses using survey data are highly valuable, they face major obstacles given that researchers do not know forecasters’ information sets, payoff functions, and the DGP. A key contribution of our paper is implementing a large-scale experiment to cleanly connect biases in expectations with both the properties of the underlying process and the forecast horizon.

Second, we contribute to experimental studies of forecasts (see [Assenza, Bao, Hommes and Massaro \(2014\)](#) for a survey). Prior work on forecasting stochastic processes includes [Hey \(1994\)](#), [Frydman and Nave \(2016\)](#) and [Beshears, Choi, Fuster, Laibson and Madrian](#)

(2013); we provide an extensive review of this literature in Table A.1. Most closely related, Reimers and Harvey (2011) also document that the forecast-implied persistence is higher than the actual persistence for transitory processes, which indicates the robustness of this phenomenon, but they do not analyze the term structure of forecasts or test models of expectations. Overall, relative to existing research, our experiments have a large scale, a wide range of settings, and diverse demographics; we also collect the term structure of forecasts. In addition, we use the data to investigate a number of commonly-used models, while prior studies tend to focus on testing a particular type of model.

Finally, we use our data to shed further light on models of expectation formation, highlighting how biases in the assessment of the long-run mean can be a parsimonious way to account for key empirical facts. Some modeling techniques we use are related to the literature on noisy perception and rational inattention (Woodford, 2003; Sims, 2003). This literature has focused on frictions in perception of new information. Instead, our model emphasizes frictions in exploiting past information, which is key for generating overreaction. Another set of models postulate that forecasters use an incorrect value of the persistence ρ (Gabaix, 2018; Angeletos, Huo and Sastry, 2021). We find that a given “mistaken” ρ cannot simultaneously account for the degree of overreaction in short-term and long-term forecasts. If using an incorrect ρ is the main bias, overreaction will dissipate for long-term forecasts, which is not the case in the data.

Several recent models examine the role of memory in belief formation, which also feature frictions in exploiting past information. Bordalo, Gennaioli and Shleifer (2020b), Bordalo, Coffman, Gennaioli, Schwerter and Shleifer (2020a), and Bordalo, Conlon, Gennaioli, Kwon and Shleifer (2023) build on representativeness (Kahneman and Tversky, 1972) and associative recall (Kahana, 2012). Wachter and Kahana (2019) present a retrieved-context theory of beliefs to model associative recall.³ The most closely related analysis in this area is da Silveira, Sung and Woodford (2020), which generalizes earlier work by da Silveira and Woodford (2019). They present a dynamic model of optimally noisy memory where past information is summarized by a memory state formed before each period; when the memory

³In addition, Nagel and Xu (2022) and Neligh (2022) study applications of memory decay. In empirical analyses, Enke, Schwerter and Zimmermann (2020) experimentally test how associative recall affects beliefs. Hartzmark, Hirshman and Imas (2021) and D’Acunto and Weber (2020) also find evidence consistent with memory playing a role in decision making.

is imprecise, the agent optimally puts more weight on the latest observation, which generates overreaction. In our model, the overweighting of recent information can be related to memory constraints, but it can also arise from other frictions in information processing that lead to heightened activation and disproportionate focus on recent information (see, e.g., [Baddeley and Hitch, 1993](#); [Spillers, Brewer and Unsworth, 2012](#); [Cowan, 2017a](#)), which we discuss in more detail in [Section 5.1](#).

Although we focus on overreaction given our empirical findings, we provide an extension of our model in [Online Appendix IA5](#) which accommodates underreaction by introducing noisy signals to the belief formation process. These noisy signals can play a role in survey data ([Coibion and Gorodnichenko, 2012, 2015](#); [Kohlhas and Walther, 2021](#)), but are unlikely to be first-order in our simple forecasting experiment. In this extension, the relative degree of overreaction is still stronger when the process is less persistent, consistent with our evidence and findings from field data in [Section 2.1](#). Meanwhile, the average level of the bias can be over- or underreaction, depending on the noisiness of signals.⁴

The rest of the paper proceeds as follows. [Section 2](#) discusses stylized facts from survey data and the limitations of these analyses, which motivate our experiment. [Section 3](#) describes the experiment. [Section 4](#) presents our main empirical facts that overreaction is stronger for less persistent processes and longer forecast horizons. It also analyzes whether commonly-used models align with the evidence. [Section 5](#) presents our model and shows that the model fits the data well. [Section 6](#) provides additional results for the mechanism of the model and performs robustness checks of model assumptions. [Section 7](#) concludes.

2 Motivating Facts

To motivate our experimental investigation, we first present stylized facts from survey forecasts and discuss their limitations.

⁴Recent experiments by [Fan, Liang and Peng \(2021\)](#) show that participants underreact in inferring, e.g., whether a company is good or bad after observing a positive stock return signal, consistent with previous results surveyed by [Benjamin \(2019\)](#). However, participants overreact in forecasting the stock return in the next period with the same signal. They find that underreaction in the inference problem may arise because the signal appears less similar to the inference problem than the forecasting one, and rephrasing the inference question to be more similar to the signal can lead to overreaction. The perceived noise of the signal in our underreaction extension can accommodate such perceived similarity between the signal and the problem.

2.1 Variations in Overreaction: Evidence from Survey Forecasts

A major challenge for analyzing expectations using survey forecast data is that the true DGP and forecasters' information sets are both unknown. Taking inspiration from [Coibion and Gorodnichenko \(2015\)](#), [Bordalo et al. \(2020c\)](#) observe that one idea is to capture belief updating using forecast revisions by individual forecasters: revisions should incorporate news that a forecaster responds to and should be part of the information set. When a forecaster overreacts to information, revisions at the individual level would overshoot (e.g., upward forecast revisions would predict realizations below forecasts). The empirical specification is the following, which regresses forecast errors on forecast revisions in a panel of quarterly individual-level forecasts:

$$\underbrace{x_{t+h} - F_{i,t}x_{t+h}}_{\text{forecast error}} = a + b \underbrace{(F_{i,t}x_{t+h} - F_{i,t-1}x_{t+h})}_{\text{forecast revision}} + v_{it}, \quad (2.1)$$

where $F_{i,t}x_{t+h}$ is the forecast of individual i of outcome x_{t+h} . For each series, we obtain a coefficient b (henceforth the "error-revision coefficient"). When overreaction is present, b should be negative, and vice versa ([Bordalo et al., 2020c](#)).

[Bordalo et al. \(2020c\)](#) analyze professional forecasts of 22 series of macroeconomic and financial variables. They find that the error-revision coefficient b is generally negative, and it is more negative for processes with lower persistence. They interpret this pattern as an indication that overreaction tends to be stronger when the actual process is more transitory. In [Figure I, Panel A](#), we use Survey of Professional Forecasters (SPF) data and replicate this finding. Here we use the simple one-period-ahead forecasts, namely $h = 1$. The y -axis shows the coefficient b for different series, and the x -axis shows the autocorrelation of each series as a simple measure of persistence. We see that the coefficient b is more negative (i.e., overreaction is stronger) when the actual series is less persistent.

In [Figure I, Panel B](#), we also document similar results using analysts' forecasts of firms' sales from the Institutional Brokers' Estimate System (IBES). Again we use one-period-ahead forecast, namely $h = 1$. We normalize both actual sales and projected sales using lagged sales, and the frequency is quarterly. Results are very similar if we use an annual frequency,

or using earnings forecasts instead of sales forecasts.⁵ We run one regression in the form of Equation (2.1) for each firm i to obtain coefficient b_i . We also compute the autocorrelation of the actual sales growth process ρ_i . Figure I, Panel B, shows a binscatter plot of the average b_i in twenty bins of ρ_i . Here the key fact remains: the coefficient b_i is more negative when the actual sales process of the firm is less persistent. Nonetheless, the average level of b_i is positive (indicating underreaction), consistent with Bouchaud et al. (2019). In contrast, Bordalo et al. (2019) find negative error-revision coefficients (pointing to overreaction) when analyzing analyst forecasts about long-term growth instead of near-term cash flows, which suggests the forecast horizon may affect the degree of overreaction as well.

2.2 Challenges in Field Data

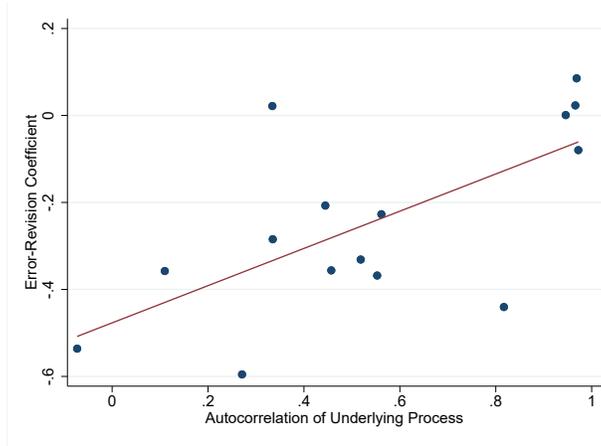
Although these results from survey data are intriguing, they can be difficult to interpret unequivocally for several reasons.

First, the error-revision regressions have limitations. To begin, it is difficult to estimate b precisely for transitory processes when expectations are close to rational. In this case, revisions are close to zero, so the regression coefficient is not well estimated. As an illustration, in Figure A.1, Panel A, we show the error-revision coefficient b from simulations where we simulate forecasters under diagnostic expectations (Bordalo et al., 2018, 2020c) for AR(1) processes with different levels of persistence. By construction, the simulated coefficient (shown by the solid line) is on average similar to theoretical predictions in the diagnostic expectations model (Bordalo et al., 2020c). Meanwhile, the dashed lines show that the confidence intervals become very wide when the process persistence is below 0.5.⁶ The intuition in this example is that the variance of the right-hand-side variable, the forecast revision, goes to zero for i.i.d. processes when expectations are close to rational (see discussion on asymp-

⁵Earnings forecasts have several complications relative to sales forecasts. First, earnings forecasts primarily take the form of earnings-per-share (EPS), which may change if firms issue/repurchase shares, or have stock splits/reverse splits. This requires us to transform EPS forecasts to implied forecasts about total firm earnings, which could introduce additional measurement error. Second, the definition of earnings firms use for EPS can be informal (“pro forma” earnings, instead of formal net income according to the Generally Accepted Accounting Principles (GAAP)). As a result, matching earnings forecasts properly with actual earnings can be more challenging. In comparison, sales forecasts are directly about total sales of the firm, and the accounting definition of sales is clear (based on GAAP).

⁶For AR(1) processes, the diagnostic forecast is $E_t^\rho x_{t+1} = E_t x_{t+1} + \rho \epsilon_t$, where $E_t x_{t+1}$ is the rational forecast, ρ is the AR(1) persistence, and ϵ_t is the news in period t . When the process is i.i.d., the diagnostic forecast becomes the same as the rational forecast, and the error-revision coefficient is not well-defined.

(A) SPF Forecasts



(B) Analyst Forecasts

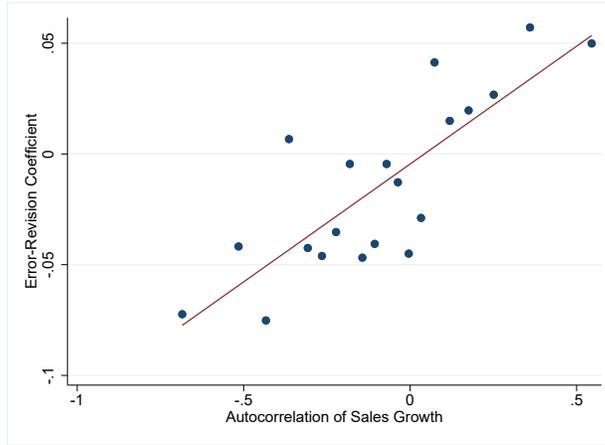


FIGURE I

Forecast Error on Forecast Revision Regression Coefficients

In Panel A, we use SPF data on macroeconomic forecasts and estimate a quarterly panel regression using each individual j 's forecasts for each variable x_i : $x_{i,t+1} - F_{i,j,t}x_{i,t+1} = a + b_i(F_{i,j,t}x_{i,t+1} - F_{i,j,t-1}x_{i,t+1}) + v_{i,j,t}$, where the left-hand-side variable is the forecast error and the right-hand-side variable is the forecast revision for each forecaster j . The y -axis plots the regression coefficient b_i for each variable, and the x -axis plots the autocorrelation of the variable. The variables include quarterly real GDP growth, nominal GDP growth, GDP price deflator inflation, CPI inflation, unemployment rate, industrial production index growth, real consumption growth, real nonresidential investment growth, real residential investment growth, real federal government spending growth, real state and local government spending growth, housing start growth, unemployment rate, 3-month Treasury yield, 10-year Treasury yield, and AAA corporate bond yield. In Panel B, we use IBES data on analyst forecasts of firms' sales and estimate a quarterly panel regression using individual analyst j 's forecasts for each firm i 's sales $x_{i,t+1} - F_{i,j,t}x_{i,t+1} = a + b_i(F_{i,j,t}x_{i,t+1} - F_{i,j,t-1}x_{i,t+1}) + v_{i,j,t}$, where the left-hand-side variable is the forecast error and the right-hand-side variable is the forecast revision for each forecaster j . The y -axis plots the regression coefficient b_i , and the x -axis plots the autocorrelation of firm i 's sales. For visualization, we group firms into twenty bins based on the persistence of their sales, and present a binscatter plot. Both actual and projected sales are normalized by lagged book assets.

otic standard errors in Online Appendix [IA2.1](#)).

Estimation issues aside, the error-revision coefficient b is not necessarily a direct metric for the degree of overreaction (i.e., how much subjective beliefs over-adjust relative to the rational benchmark). This empirical coefficient does not directly map into a structural parameter, and its interpretation can be model dependent.⁷ In addition, the error-revision coefficient b can be subject to the critique that if the forecast $F_t x_{t+h}$ is measured with noise, the regression coefficient b could be mechanically negative, given that $F_t x_{t+h}$ affects both the right side (forecast revision) and the left side (forecast error) of the regression. Ultimately, as explained earlier, one important reason for using the error-revision regression in survey forecast data is that researchers do not observe forecasters' information sets, and forecast revision can be used to capture information they respond to. If the information set is clear, however, then we can directly analyze forecasts and errors using components of the information sets.

Second, the DGP can be difficult to pin down in survey data. Many models assume the DGP to be AR(1), and the interpretation of the regression coefficients (such as Equation (2.1)) can change if the DGP is not AR(1). For instance, [Bordalo et al. \(2020c\)](#) show that the regression specification in Equation (2.1) no longer holds and a modified specification is required if the DGP is AR(2). However, econometricians may not be able to statistically differentiate whether a process is AR(1) or ARMA with longer lags in finite sample ([Fuster, Laibson and Mendel, 2010](#)).

These challenges in the field data show that complementary experimental analyses would be useful. Accordingly, we implement a large-scale forecast experiment where the forecasting environment is simple and the DGP is clearly defined, which allows us to measure over/underreaction precisely. The experiments also allow us to randomly assign participants into different conditions, so we can cleanly document the properties of subjective forecasts in different settings.

In the experiment, for AR(1) processes a straightforward way to measure the degree of over/underreaction is to examine the sensitivity of subjective forecasts to realized observa-

⁷In particular, since the forecast revision in period t is the change between the subjective forecast from $t - 1$ to t ($F_t x_{t+h} - F_{t-1} x_{t+h}$), its size and variance are affected by the past forecast ($F_{t-1} x_{t+h}$), so the magnitude of the error-revision coefficient b can be path dependent.

tions. We can regress the forecast $F_t x_{t+h}$ on x_t to obtain the forecast-implied persistence for horizon h , $\rho_{s,h}$, and then calculate $\rho_h^s = \rho_{s,h}^{1/h}$ as the implied one-period persistence of the horizon h forecasts. We can compare the implied persistence ρ_h^s with the actual persistence ρ of the process. When $\rho_h^s > \rho$, there is overreaction, in the sense that the forecast displays excess sensitivity to the latest observation x_t (i.e., when x_t is high, the forecast tends to be too high, and vice versa). Relative to the error-revision coefficient, the magnitude of the implied persistence (ρ_h^s) is easier to interpret as a measure of overreaction: the gap between ρ_h^s and ρ directly captures how much more forecasts respond to recent observations relative to the rational benchmark.⁸ In addition, this approach does not face the econometric complications that apply to the error-revision coefficient we discussed above, as the variance of the right-hand-side variable (namely the past realization) is always well-defined (it does not vanish to zero as ρ decreases). Figure A.1, Panel B, presents simulations of the implied persistence (with $h = 1$) and shows this approach is reliable for all levels of persistence.

We explain the experimental design in the next section. Our experiments focus on AR(1) processes for three reasons based on the previous discussion. First, the rational benchmark is clear in this setting. Second, we can precisely measure the degree of over-/under-reaction relative to the benchmark. Third, most models of expectations formation also focus on AR(1), so we can compare our findings with their predictions.

3 Experiment Design

We design a simple forecasting experiment, where the DGP is an AR(1) process:

$$x_{t+1} = (1 - \rho)\mu + \rho x_t + \epsilon_t. \quad (3.1)$$

The experiment begins with a consent form, followed by instructions and tests. Participants first observe 40 past realizations of the process. Then, in each round, participants make fore-

⁸There is an approximate relationship between the error-revision coefficient and $\zeta(\rho, h) = (\rho_h^s)^h / \rho^h$. Specifically, $1/(1+b) = \frac{\text{Var}(FR)}{\text{Cov}(FE+FR, FR)}$, where FR and FE refer to forecast revision and forecast error, respectively. If we set $F_{t-1}x_{t+h}$ as a constant, then this coefficient is the same as $\zeta(\rho, h)$. Accordingly, a negative error-revision coefficient, often interpreted as evidence of overreaction, implies $\zeta(\rho, h) > 1$, i.e., overreaction of the subjective belief to the latest observation.

casts and observe the next realization, for 40 rounds. After the prediction task, participants answer some basic demographic questions.

Each participant is only allowed to participate once. Participants include both individuals across the US from Amazon’s online Mechanical Turk platform (MTurk) and MIT undergraduates in Electrical Engineering and Computer Science (EECS). For MTurk, we use HITs titled “Making Statistical Forecasts.”⁹ For MIT students, we send recruiting emails to all students with a link to the experimental interface.

3.1 Experimental Conditions

There are four main sets of experiments, which we describe below and summarize in Table A.2 in the Online Appendix.

Experiment 1 (Baseline, MTurk). Experiment 1 is our baseline test, conducted in February 2017 on MTurk. We use six values of ρ : $\{0, 0.2, 0.4, 0.6, 0.8, 1\}$. The volatility of ϵ is 20. The constant μ is zero. Participants are randomly assigned to one value of ρ . Each participant sees a different realization of the process. At the beginning, participants are told that the process is a “stable random process.” In each round, after observing realization x_t , participants predict the value of the next two realizations $F_t x_{t+1}$ and $F_t x_{t+2}$. The previous period forecast of x_{t+1} , $F_{t-1} x_{t+1}$, is reported as a grey dot (so that the forecaster remembers the past forecast).¹⁰ Figure A.2 provides a screenshot of the prediction page. There are 207 participants in total and about 30 participants per value of ρ .

Experiment 2 (Long Horizon, MTurk). Experiment 2 investigates longer horizon forecasts. We assign participants to conditions identical to Experiment 1, except that we collect forecasts of x_{t+1} and x_{t+5} (instead of x_{t+2}), with $\rho \in \{0.2, 0.4, 0.6, 0.8\}$. Experiment 2 was conducted in June 2017 on MTurk. There are 128 participants in total.

⁹The MTurk platform is commonly used in experimental studies (Kuziemko, Norton, Saez and Stantcheva, 2015; D’Acunto, 2020; Cavallo, Cruces and Perez-Truglia, 2017; DellaVigna and Pope, 2017, 2018). It offers a large subject pool and a more diverse sample compared to lab experiments. Prior research also finds the response quality on MTurk to be similar to other samples and to lab experiments (Casler, Bickel and Hackett, 2013; Lian, Ma and Wang, 2018).

¹⁰This feature is present in most of our conditions, but it has no effect on our core results. In Online Appendix Table A.4, we report the forecast-implied persistence (regression of $F_t x_{t+1}$ on x_t) for conditions with and without the grey dot. We see no difference in the results among these conditions.

Experiment 3 (Describe DGP, MIT EECS). In Experiment 3, we study whether providing more information about the DGP affects forecasts. To make sure that participants have a good understanding of the AR(1) formulation, we perform this test among MIT undergraduates in Electrical Engineering and Computer Science (EECS). Experiment 3 was conducted in March 2018 and there are 204 participants. We use the same structure as in Experiment 1, with AR(1) persistence $\rho \in \{0.2, 0.6\}$. For each persistence level, the control group is the same as Experiment 1, and the process is described as “a stable random process.” For the treatment group, we describe the process as “a fixed and stationary AR(1) process: $x_t = \mu + \rho x_{t-1} + e_t$, with a given μ , a given ρ in the range $[0,1]$, and e_t is an i.i.d. random shock.” Thus there are $2 \times 2 = 4$ conditions in total, and participants are randomly allocated to one of them. At the end of the experiment, we further ask students questions testing their prior knowledge of AR(1) processes. We do not disclose the values of μ and ρ , since the objective of our study is to understand how people form forecasting rules; directly providing the values of μ and ρ would make this test redundant.

Experiment 4 and 5 (Additional Test, MTurk). In Experiment 4 and 5, we perform additional tests explained in more detail in Section 6.1 and Section 6.2. Participants are randomly assigned into a given condition and a given level of ρ . As before, there are about 30 participants for each treatment condition and level of ρ . Experiment 4 was conducted in March 2021 on MTurk and Experiment 5 was conducted in June 2022 on MTurk.

3.2 Payments

We provide fixed participation payments as well as incentive payments that depend on the performance in the prediction task. For the incentive payments, participants receive a score for each prediction that increases with the accuracy of the forecast (Dwyer, Williams, Battalio and Mason, 1993; Hey, 1994): $S = 100 \times \max(0, 1 - |\Delta|/\sigma)$, where Δ is the difference between the prediction and the realization, and σ is the volatility of the noise term ϵ . This loss function ensures that a rational participant will optimally choose the rational expectation, and it ensures that payments are always non-negative. A rational agent would expect

to earn a total score of about 2,800.¹¹ We calculate the cumulative score of each participant, and convert it to dollars. The total score is displayed on the top left corner of the prediction screen.

For experiments on MTurk (Experiments 1, 2, 4, and 5), the base payment is \$1; the conversion ratio from the score to dollars is 600, which translates to incentive payments of about \$5 for rational agents. For experiments with MIT students (Experiment 3), the base payment is \$5; the conversion ratio from the score to dollars is 240, which translates to incentive payments of about \$12 for rational agents.

3.3 Summary Statistics

Table A.3 shows participant demographics and other experimental statistics. Overall, MTurk participants are younger and more educated than the U.S. population. The mean duration of the experiment is about 18 minutes, and the hourly compensation is in the upper range of tasks on MTurk. As expected, MIT EECS undergrads are younger. Their forecast duration and overall forecast scores are similar to the MTurk participants.

4 Main Empirical Findings

In this section, we present the main findings from our experiments. In Section 4.1, we present the key empirical facts. In Section 4.2, we then analyze whether commonly-used models of expectations are in line with these facts.

As mentioned in Section 2, the error-revision regression approach has limitations for accurately measuring the degree of overreaction. In our experiment, a natural and more precise alternative measure of the degree of overreaction is the persistence implied by the forecast. Specifically, denote $\rho_{s,h}$ as the coefficient in the regression:

$$F_{it}x_{t+h} = c + \rho_{s,h}x_t + u_{it}, \quad (4.1)$$

¹¹ $E(1 - |x_{t+1} - F_t|/\sigma)$ is maximal for a forecast F_t equal to the 50th percentile of the distribution of x_{t+1} conditional on x_t . Given that our process is symmetric around the rational forecast, the median is equal to the mean, and the optimal forecast is equal to the conditional expectation. Whether the rational agent knows the true ρ (Full Information Rational Expectations) or predicts realizations using linear regressions (Least-Square Learning) does not change the expected score by much. In simulations, over 1,000 realizations of the process, we find that expected scores of the two approaches differ by less than 0.3%.

for each level of AR(1) persistence ρ and forecast horizon h (estimated in the panel of individual-level forecasts); we can also run the regression for every individual forecaster and take the mean or median regression coefficient for each process persistence ρ and forecast horizon h , which produces similar results. The implied persistence is then given by $\rho_h^s = \rho_{s,h}^{1/h}$. As the Full Information Rational Expectation (FIRE) is given by $\rho^h x_t$, the difference between ρ_h^s and ρ provides a direct measure of the degree of overreaction. This measure is reliable for AR(1) processes as we show in Section 2, and forecasters' information sets are clear in the experiment.

4.1 Basic Facts: More Overreaction for More Transitory Processes and Longer Forecast Horizons

A. Overreaction and Process Persistence

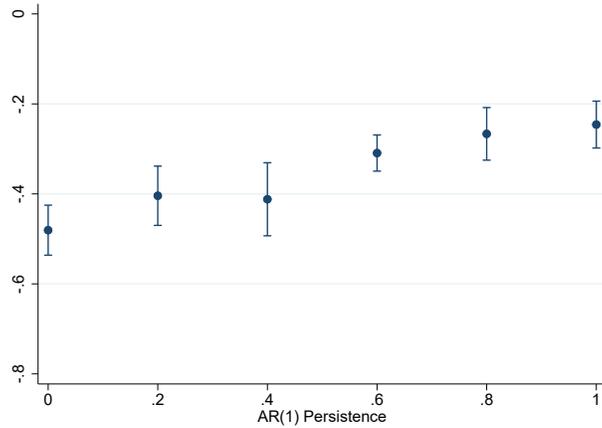
We first describe how overreaction varies for AR(1) processes with different levels of persistence, starting with one-period-ahead forecasts. Using data from Experiment 1, we have AR(1) processes with persistence from 0 to 1. First, in Figure II, Panel A, we run the error-revision regression in Equation (2.1), as we did using field data (the y -axis shows the error-revision coefficient, and the x -axis shows the persistence of the process). We find that the coefficient b is more negative for transitory processes, in line with the evidence from field data discussed in Section 2. Second, in Figure II, Panel B, we plot the implied persistence ρ_1^s against the actual process persistence ρ . We see that when $\rho = 1$, ρ_1^s is around one (i.e., the subjective and rational forecasts have roughly the same sensitivity to x_t). When ρ is smaller, ρ_1^s declines, but not as much. When $\rho = 0$, ρ_1^s is around 0.45 (i.e., the sensitivity of the subjective forecast to x_t is much larger than that under the rational benchmark).¹²

Overall, in the experiment, by explicitly controlling the DGP and forecasters' information sets, we can establish clearly that overreaction is stronger for more transitory processes. We discuss several robustness checks below.

FIRE vs. In-Sample Least Square Learning. The analyses above use FIRE for the actual

¹²We can also compute the ratio of relative overreaction $\zeta(\rho, h) = \frac{\rho_{s,h}^s}{\rho^h}$. Figure A.3 plots the value of $\zeta(\rho, h)$ for each level of ρ (except when $\rho = 0$ where $\zeta(\rho, h)$ is not well-defined). Since ρ_1^s decreases less than one-for-one with ρ , the degree of overreaction is higher when the process is less persistent.

(A) Forecast Error on Forecast Revision Regression Coefficients



(B) Forecast-Implied Persistence and Actual Persistence

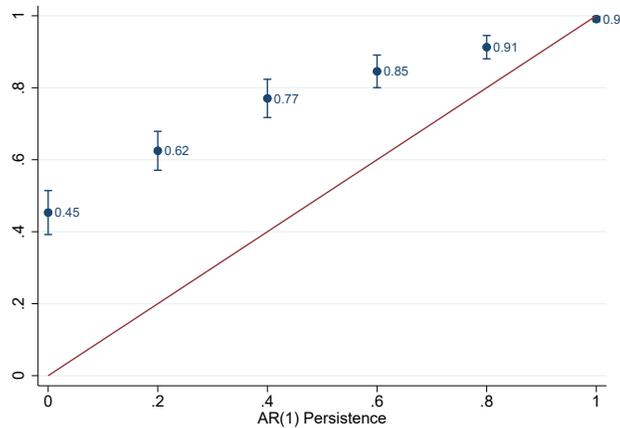


FIGURE II

Implied Persistence and Actual Process Persistence

In Panel A, we use data from Experiment 1 and for each level of AR(1) persistence ρ , we estimate a panel regression of forecast errors on forecast revisions: $x_{t+1} - F_{i,t}x_{t+1} = a + b(F_{i,t}x_{t+1} - F_{i,t-1}x_{t+1}) + v_{it}$. The y -axis plots the regression coefficient b , and the x -axis plots the AR(1) persistence ρ . In Panel B, we estimate the implied persistence ρ_1^s from $F_{it}x_{t+1} = c + \rho_1^s x_t + u_{it}$ for each level of AR(1) persistence ρ . The y -axis plots the implied persistence ρ_1^s , and the x -axis plots the AR(1) persistence ρ . The red line is the 45-degree line, and corresponds to the implied persistence under Full Information Rational Expectations (FIRE). The vertical bars show the 95% confidence interval of the point estimates.

persistence ρ . Results are similar if we use in-sample least square learning as the rational benchmark instead. Specifically, the in-sample least square estimates are formed as:

$$\widehat{E}_t x_{t+h} = \widehat{a}_{t,h} + \sum_{k=0}^{k=n} \widehat{b}_{k,h,t} x_{t-k}. \quad (4.2)$$

In period t the forecaster predicts x_{t+h} using lagged values from x_{t-k} up to x_t ; parameters $\widehat{a}_{t,h}$ and $\widehat{b}_{k,h,t}$ are estimated, on a rolling basis, using OLS and past realizations until x_t . The estimated coefficients may differ based on persistence ρ . We set $n = 3$, but results are not sensitive to the number of lags.

In our data, the difference between $\widehat{E}_t x_{t+h}$ and FIRE is small. The top panel of Figure A.4 shows that the mean squared difference between these two expectations is small, and does not decrease much after 40 periods. This is because our AR(1) processes are very simple, and a few dozen data points are enough for least square forecasts to be reasonably accurate. It also shows that the mean squared difference between the least square forecast and the actual forecasts is substantial, and does not change much across different periods. The bottom panel shows that the persistence implied by least square learning is about the same as the actual persistence ρ . Accordingly, in the rest of the paper we use FIRE in our baseline definitions, but all the results are similar if we use the in-sample least square $\widehat{E}_t x_{t+h}$ instead.

Effect of Linear AR(1) Prior. We also analyze whether explicitly providing a prior that the DGP is a linear AR(1) process affects the results. In Experiment 1 with participants from the general population, we describe the process as a “stable random process” (given that most of these participants may not know what an AR(1) process means). In Experiment 3 with MIT EECS students, we tell half of the participants that the DGP is AR(1) with fixed μ and ρ (treatment group), and half of the participants the process is a “stable random process” (control group). Given the size limit of MIT EECS students, we use two values of ρ : 0.2 and 0.6 (so participants are randomly assigned into a given value of ρ and a given type of process description).

In Figure A.5, we show that whether the linear AR(1) information was provided has no discernible impact on the properties of forecast errors. In Panel A, we plot the distributions of the forecast errors, which are almost identical in the treatment vs. control group. In Panel

B, we find that the predictability of forecast errors conditional on the latest observation x_t is also similar in the treatment vs. control group. In both samples, forecasts tend to be too high when x_t is high (overreaction), and the magnitude of the bias is about the same. Table A.5 shows that the implied persistence is also similar in both the treatment and control groups. Overall, we find that explicit descriptions of the AR(1) process do not seem to affect the basic patterns in the data. Put differently, participants do not seem to enter the experiment with complicated nonlinear priors. Finally, Figure A.6 compares the implied persistence ρ_1^s in the MIT experiments with that in our baseline experiments, which shows the results are very stable (even with the caveat that these two sets of experiments were not conducted and randomized at the same time).

Stability across Demographics. Figure A.7 shows both the error-revision coefficient b and implied persistence ρ_1^s against ρ in different demographic groups. In all cases, the main patterns are stable.

B. Overreaction and Forecast Horizon

We next examine how overreaction varies across different forecast horizons, by comparing forecasts for x_{t+1} and forecasts for x_{t+2} , x_{t+5} , and x_{t+10} . Recent research using survey data suggests that overreaction appears more pronounced for forecasts of longer-horizon outcomes. For instance, using the error-revision regression, Bordalo et al. (2019) find a negative and significant coefficient for equity analysts' forecasts of long-term earnings growth (which points to overreaction), while Bouchaud et al. (2019) document a positive error-revision coefficient for analysts' forecasts of short-term earnings (which points to underreaction). Based on professional forecasters' predictions of interest rates, several studies also show that the error-revision coefficient is negative and significant for long-term interest rates, but not for short-term interest rates (Bordalo et al., 2020c; Wang, 2021; d'Arienzo, 2020).

We compare the degree of overreaction for long-horizon versus short-horizon forecasts following the structure in Giglio and Kelly (2018). They show that affine asset pricing models using a given level of process persistence cannot simultaneously account for prices of long-maturity and short-maturity claims. We ask if the level of implied persistence (ρ_h^s) differs for long-term and short-term forecasts in our data.

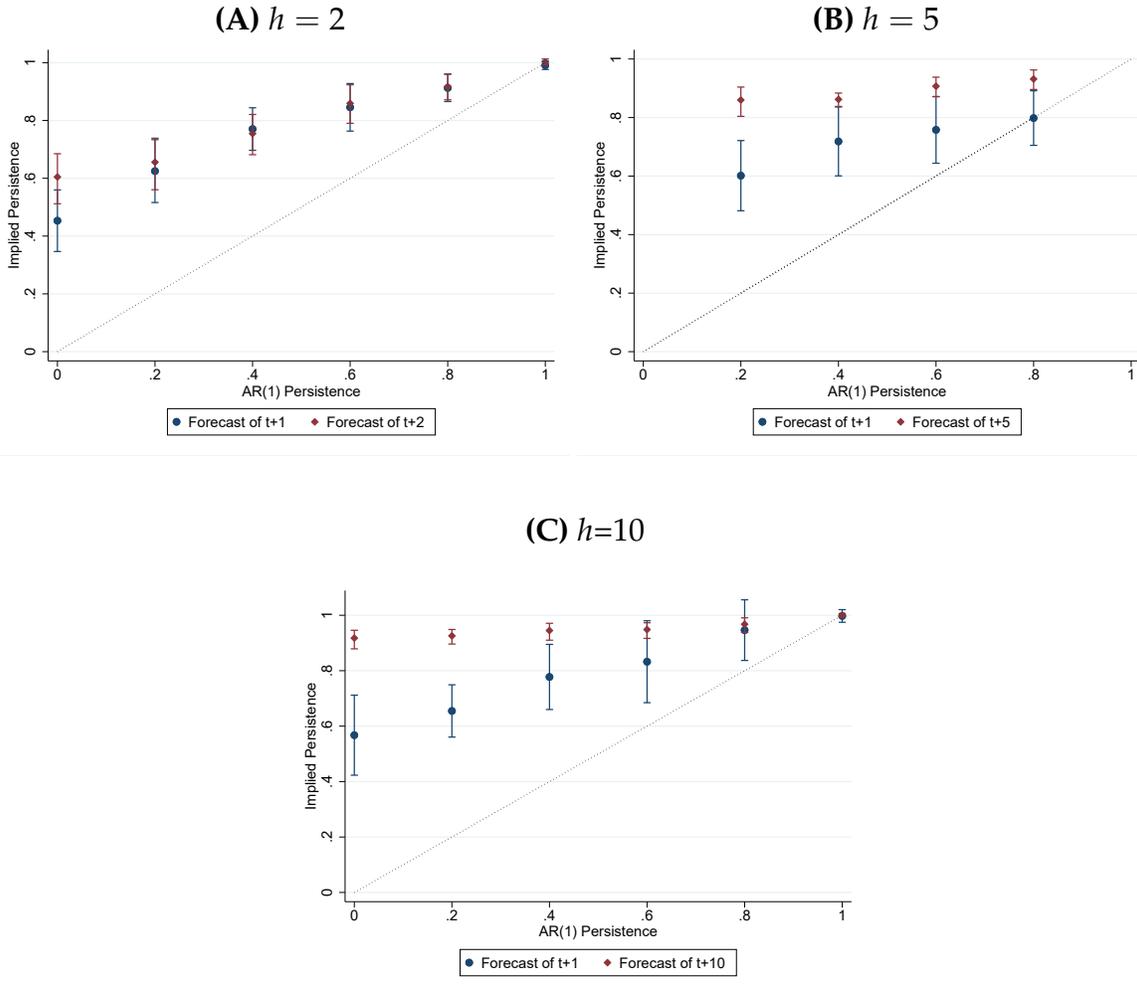


FIGURE III

Implied Persistence for Short-Term and Long-Term Forecasts

This figure shows the implied persistence ρ_h^s as a function of the actual persistence ρ . The implied persistence ρ_h^s is obtained by regressing $F_t x_{t+h}$ on x_t and taking the $1/h$ th power of the coefficient. The vertical bars indicate the 95% confidence interval. Panels A, B, and C show results for $h = 2$, $h = 5$, $h = 10$, respectively. The data for Panels A, B, and C come from Experiment 1 (where participants forecast x_{t+1} and x_{t+2}), 2 (where participants forecast x_{t+1} and x_{t+5}), and 4 (where participants forecast x_{t+1} and x_{t+10}) respectively. The two sets of dots (short-horizon and longer-horizon forecasts) for each panel come from forecasts made by the same participants. The dotted line is the 45-degree line. There are no results for $\rho = 0$ and $\rho = 1$, because the experiment initially designed to collect data on $h = 5$ (Experiment 2) did not include these conditions.

Figure III shows the results. Panels A, B, and C report the values of the implied persistence ρ_h^s for $h = 2$, $h = 5$ (for which we only have conditions with ρ between 0.2 and 0.8), and

$h = 10$, respectively. In each panel, the longer-horizon forecasts and the short-horizon forecasts (for x_{t+1}) come from the same participants for the same AR(1) process. In all panels, we see a substantial degree of overreaction. Moreover, the degree of overreaction (as reflected by the gap between the implied persistence ρ_h^s and ρ) is even higher when h is larger. The difference between the implied persistence for short-horizon and longer-horizon forecasts is especially pronounced for very large h and for transitory processes (significant at 5% as the confidence intervals show). The results also suggest that forecasters using a single incorrect persistence parameter (e.g., [Gabaix, 2018](#)) cannot fully explain the empirical evidence: this bias alone cannot simultaneously square with both short-term and long-term forecasts made by the same forecasters.

4.2 Testing Models of Expectations

We now use the data from our experiments and the key facts above to analyze commonly-used models of expectations formation.

A. Models of Expectations

We begin by laying out the standard formulations of these expectations models.

Backward-Looking Models

We start with older “backward-looking” models, which specify fixed forecasting rules based on past data and do not incorporate properties of the process (e.g., they are not a function of ρ). The term structure of expectations in these models is not well-defined, so we focus on one-period-ahead forecasts.

1. Adaptive expectations

Adaptive expectations have been used since at least the work of [Cagan \(1956\)](#) on inflation and [Nerlove \(1958\)](#) on cobweb dynamics. The standard specification is:

$$F_t x_{t+1} = \delta x_t + (1 - \delta) F_{t-1} x_t. \tag{4.3}$$

2. Extrapolative expectations

Extrapolative expectations have been used since at least [Metzler \(1941\)](#), and are sometimes used in studies of financial markets ([Barberis, Greenwood, Jin and Shleifer, 2015](#); [Hirshleifer, Li and Yu, 2015](#)). One way to specify extrapolation is:

$$F_t x_{t+1} = x_t + \phi(x_t - x_{t-1}). \quad (4.4)$$

That is, expectations are influenced by the current outcome and the recent trend, and $\phi > 0$ captures the degree of extrapolation.

Forward-Looking Models

We now proceed to “forward-looking” models, where forecasters do incorporate features of the true process. In these models, which contain rational expectations, the term structure of expectations is well-defined.

3. Full information rational expectations

Full information rational expectations (FIRE) is the benchmark specification in economic modeling. Decision makers know the true DGP and its parameters, and make statistically optimal forecasts accordingly:

$$F_t x_{t+h} = E_t x_{t+h} = \rho^h x_t. \quad (4.5)$$

As explained in [Section 4.1](#), in our data in-sample least square learning is very close to FIRE, so we use FIRE as the benchmark.

4. Noisy information/sticky expectations

Noisy information models assume that forecasters do not observe the true underlying process, but only noisy signals of it (e.g., [Woodford, 2003](#)). In our setting where actual realizations are displayed directly, such frictions may correspond to noisy perception. These models typically have the following recursive formulation:

$$F_t x_{t+h} = (1 - \lambda) E_t x_{t+h} + \lambda F_{t-1} x_{t+h} + \epsilon_{it,h}, \quad (4.6)$$

where $E_t x_{t+h}$ is FIRE, $\lambda \in [0, 1]$ depends on the signal’s noisiness, and $\epsilon_{it,h}$ comes from the

noise too. As a reduced form formulation, Equation (4.6) can also represent anchoring on past forecasts, and [Bouchaud et al. \(2019\)](#) use it to model forecasts of equity analysts.

5. Diagnostic expectations

Diagnostic expectations are introduced by [Bordalo, Gennaioli and Shleifer \(2018\)](#) to capture overreaction in expectations driven by the representativeness heuristic ([Kahneman and Tversky, 1972](#)). The baseline specification is:

$$F_t x_{t+h} = E_t x_{t+h} + \theta(E_t x_{t+h} - E_{t-1} x_{t+h}). \quad (4.7)$$

That is, the subjective expectation is the rational expectation plus the surprise (measured as the change in rational expectations from the past period) weighted by θ , which indexes the severity of the bias. Under diagnostic expectations, subjective beliefs adjust to the true process and incorporate features of rational expectations ("kernel of truth"), but overreact to the latest surprise by degree θ .

6. Constant gain learning

We also test constant gain learning about x_t , which implies least square learning with weights that decrease for observations further in the past. We use the regression specification:

$$F_t x_{t+h} = \widehat{E}_t^m x_{t+h} = \widehat{a}_{h,t} + \widehat{b}_{h,t} x_t, \quad (4.8)$$

where $\widehat{a}_{h,t}, \widehat{b}_{h,t}$ are obtained through a rolling regression with all data available until t . The difference with the standard least square learning specification is that this regression uses decreasing weights (i.e., older observations receive a lower weight) to reflect imperfect retention of past information. Specifically, in period t , for all past observations $s \leq t$, we use exponentially decreasing weights: $w_t^s = \frac{1}{\kappa^{(t-s)}}$. These weights correspond to constant gain learning in recursive least squares formulations ([Malmendier and Nagel, 2016](#)).

Other Models

In addition to the models listed above, there are several other models to consider. We do not estimate these models formally because their features are qualitatively different from

the results observed in our data, by design or by outcome, as we explain below.

First, an intuitive model of overreaction is described in [Gabaix \(2018\)](#). Specifically, the forecaster faces a range of possible processes with varying degrees of persistence. To limit computational cost, the boundedly rational forecaster uses a persistence parameter $\hat{\rho}$ that is anchored to a default level of persistence ρ^d : $\hat{\rho} = m\rho_i + (1 - m)\rho^d$. In such a setting, forecasters would overreact to processes that are less persistent than the default level ρ^d , and underreact to processes that are more persistent than the default level. One limitation of this approach is such a bias alone cannot account for results across different forecast horizons. As [Figure III](#) shows, a given level of incorrect persistence cannot simultaneously square with both short-term and long-term forecasts made by the same forecasters. Indeed, if a single incorrect ρ was the only bias, then overreaction would dissipate for long-term forecasts (e.g., forecast of x_{t+10}), which is not the case in the data.

Second, several papers investigate belief formation with model misspecification. For instance, in natural expectations ([Fuster, Laibson and Mendel, 2010](#)), the key observation is that forecasters can have difficulty differentiating processes with hump-shaped dynamics (such as AR(2) or ARMA(p,q)) from simpler AR(1) processes in finite samples.¹³ Other models analyze subjective beliefs about regime shifts ([Barberis, Shleifer and Vishny, 1998](#); [Bloomfield and Hales, 2002](#); [Massey and Wu, 2005](#)). As explained in [Section 4.1](#), in Experiment 3 among MIT EECS students, we explicitly describe the linear AR(1) process to half of the participants. We do not find that the information of a linear AR(1) prior affects the results. Indeed, our findings highlight that systematic biases in expectations can be significant even in linear stationary environments.

Relatedly, building on [Rabin \(2002\)](#), [Rabin and Vayanos \(2010\)](#) also formulate a model based on beliefs about misspecified DGP. Proposition 6 in their paper states that the forecast should have a negative loading on the most recent observation (x_{t-1}), whereas this loading is strongly positive in our data.

¹³[Fuster, Laibson and Mendel \(2010\)](#) formulate an “intuitive model” $F_t x_{t+1} = x_t + \phi(x_t - x_{t-1}) + \epsilon_{t+1}$, when the true DGP is an AR(2) $x_{t+1} = \alpha x_t + \beta x_{t-1} + \eta_{t+1}$, and $\phi = (\alpha - \beta - 1)/2$. We could test this model in our data, where $\alpha \geq 0, \beta = 0, \phi < 0$, and the intuitive model has the same functional form as the extrapolative expectation in [Equation \(4.4\)](#) with negative ϕ .

B. Estimating Models of Expectations

We now estimate the models described above on one-period-ahead forecast data (i.e., with $h = 1$). We pool data from all conditions of Experiment 1 (i.e., with $\rho \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$). All models except FIRE (which has no parameter) and constant gain learning (whose parameter lies in the decreasing weights) can be simply estimated using constrained least squares. We cluster standard errors at the individual level. The constant gain learning model is estimated by minimizing, over the decay parameter, the mean squared deviation between model-generated and observed forecasts. We estimate standard errors for this model by block-bootstrapping at the individual level.

Table A.6 reports the estimated parameters. Each model is described by an equation and a parameter (in bold). The parameter estimate is reported in the third column, along with standard errors in the fourth column. In the fifth column, we report the mean squared error of each model, as a fraction of the sample variance of forecast. Since forecasts in the $\rho = 1$ condition are mechanically more variable than forecasts in the $\rho = 0$ condition, we compute one such ratio per level of ρ , and then compute the average ratio across values of ρ .

Several patterns emerge from the model estimation. First, consistent with findings in Section 4.1, rational expectations are rejected. Indeed, rational expectations are nested in all three forward-looking non-RE models, and the coefficient related to deviations from rational expectations is always significant at 1%. In line with such deviations being important, FIRE has the lowest explanatory power of forecast data.

Second, most models point to strong signs of overreaction. The adaptive model features overreaction through the high loading on the past realization x_t (0.83). The backward-looking extrapolative model has a negative coefficient on the slope ($x_t - x_{t-1}$), but this again reflects that most overreaction is built into the coefficient on the past realization x_t , which is fixed at one by definition. The diagnostic expectations model has $\theta = 0.34$, which indicates strong overreaction (forecasts react 34% “too much” to the last innovation).¹⁴ The constant gain learning model features a significant decay in the weight of past observations, a loss of 6% per period (i.e., it takes about 12 periods to divide the weight by 2), rejecting the equal

¹⁴The θ estimate is slightly lower than the typical estimate in [Bordalo et al. \(2020c\)](#) using macro survey data (which find θ of around 0.5) and in [Bordalo, Gennaioli and Shleifer \(2018\)](#) and [Bordalo et al. \(2019\)](#) using analyst forecasts of credit spreads and long-term EPS growth (which find θ of around 1).

weights in benchmark least square learning. Last, the sticky/noisy expectations model is the only one that does not feature overreaction. The coefficient on previous forecasts ($F_{t-1}x_{t+1}$) is 0.14, similar to earlier analyses on earnings forecasts by stock analysts (Bouchaud et al., 2019). This finding suggests that there is some anchoring on the level of past forecasts, in addition to overreaction to the recent realization.

C. Comparing Model Predictions and Empirical Results

We now investigate how overreaction varies with process persistence and forecast horizon in the estimated models, and compare model predictions with our empirical findings.

Overreaction and Process Persistence. We start with the evidence on overreaction and process persistence (using one-period-ahead forecasts). In Figure IV, we compute the implied persistence based on the five models estimated above. Specifically, for each model m and for each observation in our data, we compute the model-based forecast $\widehat{F}_t^m x_{t+1}$ using the parameters in Table A.6. We then group observations per level of $\rho \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$. For each level, we regress the model-based forecast $\widehat{F}_t^m x_{t+1}$ on x_t to obtain the implied persistence according to the model.

In Figure IV, the solid line represents the implied persistence based on actual forecasts (same as Figure II, Panel B). The dots represent the implied persistence based on the models. In all models, the implied persistence is an increasing function of ρ , and is close to one for random walks (as in rational expectations). However, the commonly-used models do not match the transitory processes very well. The backward-looking expectations models generate too much overreaction for transitory processes, whereas most of the forward-looking models do not generate enough overreaction. Sticky expectations generate no overreaction. The baseline specification of diagnostic expectations generates overreaction, though not for i.i.d. processes (e.g., the rational surprise in the benchmark formulation of diagnostic expectations is zero in this case). The constant gain learning model also generates some excess sensitivity to recent realizations by giving them larger weights, but the weights on past observations do not decrease fast enough.

To connect with results in field data and for completeness, we also report in Figure A.8 the error-revision coefficients based on the models. Again, the solid line represents exper-

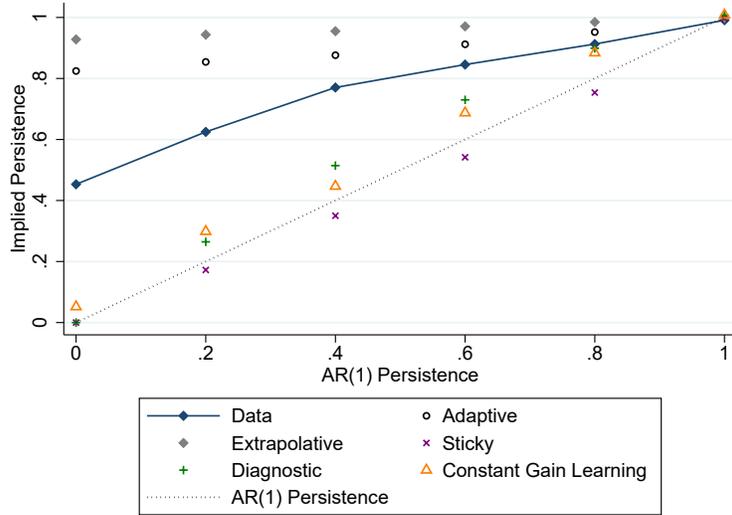


FIGURE IV
Forecast-Implied Persistence: Data vs Models

For each model m , we compute the model-based forecast $\widehat{F}_t^m x_{t+1}$ for each observation in our data. We use the model parameters reported in Table A.6. We then group observations per level of actual persistence $\rho \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$. For each level of ρ , we regress the model-based forecast $\widehat{F}_t^m x_{t+1}$ on lagged realization x_t . The dots report this regression coefficient, which is the forecast-implied persistence according to model m for a given level of ρ . The solid line corresponds to the forecast-implied persistence in the data, also shown in Figure II, Panel B.

imental data and the dots represent predictions from estimated models. Here we omit the adaptive and extrapolative models, because they do not impose an obvious structure on the two-period-ahead forecasts $F_t x_{t+2}$, which are needed to compute revisions. The conclusions are similar to those in Figure IV. For transitory processes, sticky expectations lead to error-revision coefficients that are too high (no overreaction). The baseline specification of diagnostic expectations performs well for more persistent processes, and less so for transitory processes. Constant gain learning, meanwhile, generates a coefficient that is too negative for i.i.d. processes.¹⁵

Overreaction and Forecast Horizon. We now compare the data and model predictions for longer-horizon forecasts. We focus on the forward-looking models as the backward-

¹⁵This is in fact a mechanical effect of the error-revision coefficient, which divides by the variance of forecast revision. In the constant gain learning model, forecast revisions tend to be very small for low values of ρ s (they are close to zero), which blows up the absolute value of the error-revision coefficient. The implied persistence measure in Figure IV is immune to this problem.

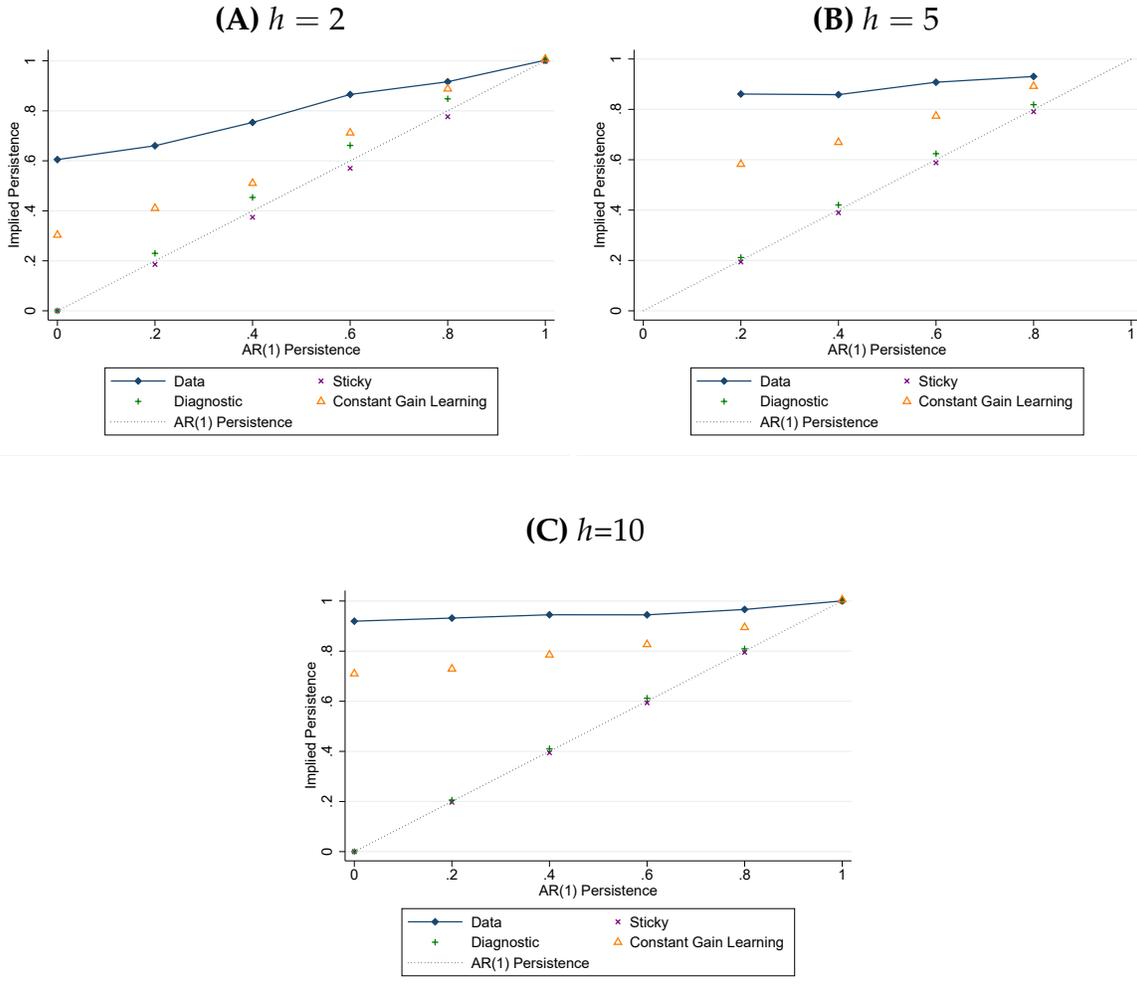


FIGURE V

Forecast-Implied Persistence: Longer Horizon Forecasts

This figure shows the implied persistence ρ_h^s as a function of the actual persistence ρ . The implied persistence ρ_h^s is obtained by regressing $F_t x_{t+h}$ on x_t and taking the $1/h$ th power of the coefficient. Panels A, B, and C show results for $h = 2, h = 5, h = 10$, respectively. The data is the same as those used in Figure III. The dotted line is the 45-degree line.

looking models do not provide a clear term structure of forecasts for multiple horizons. In particular, we fit all models using $h = 1$ and use the same parameters to generate model predictions for $h = 2, 5, 10$. We find that the implied persistence according to standard models tends to be too low, especially when the forecast horizon is longer; the exception is constant gain learning, which produces a closer fit when the horizon is longer.

Overall, the empirical findings suggest that forecasts in the data do adapt to the setting

(with lower implied persistence when the true persistence ρ is smaller). Consequently, the backward-looking models (which specify a fixed dependence of forecasts on past realizations) generally do not perform well in capturing the degree of overreaction across different levels of process persistence. Meanwhile, despite the partial adaptation of forecasts to the actual persistence, overreaction is most significant when the process is transitory. This result echoes strong extrapolative beliefs observed in other settings with transitory processes, including experimental studies (Frydman and Nave, 2016) and survey data on stock returns (Greenwood and Shleifer, 2014). The baseline specifications of the forward-looking models generally do not allow transitory shocks to generate much overreaction. This is because these models focus on using the recent shock ϵ_t to infer future shocks ϵ_{t+h} , which tends to produce an inferred value ($F_t[\epsilon_{t+h}|\epsilon_t]$) that scales with the persistence ρ .

The empirical evidence points to the following insight: a parsimonious way for models to account for the empirical findings is allowing recent observations to influence the assessment of more persistent aspects of the process (e.g., its long-run mean). The benefit of modeling imperfect inference about the long-run mean is two-fold. First, this approach naturally implies that recent observations can lead to biased forecasts of future outcomes even when the process is i.i.d. Second, this approach also naturally leads to greater overreaction of longer horizon forecasts, which are more sensitive to biases regarding the long-run mean. For the models discussed above, for example, one can modify constant gain learning to focus on the mean instead of x_{t+h} , as in Nagel and Xu (2022). Similarly, da Silveira, Sung and Woodford (2020) provide a model where the noisiness of the memory state leads the agent to put a large weight on the recent observation when inferring the mean. The biases about the long-run mean can be generalized to other models as well. As an example, for the diagnostic expectations model of Bordalo et al. (2018), a modification could be that the agent starts with a default prior regarding μ at a given point in time, and a high realization of x_t is diagnostic of a higher μ .¹⁶

In the next section, we present a simple and tractable model that builds on this insight of biases about the mean. The model delivers three key features. First, it microfound the overweighing of recent observations as a result of imperfect information processing (Sections

¹⁶This approach is similar to that of Bordalo et al. (2019), where agents partially conflate transitory noise with a more persistent underlying latent process. (In our setting, however, there is no unobservable state.)

5.1 and 5.2). Second, it can be easily estimated to evaluate empirical performance and readily generalized to different processes and settings (Section 5.4 and Online Appendix IA4). Third, it can accommodate both overreaction and underreaction (Online Appendix IA5).

5 Model

In this section, we provide a simple model that captures the disproportionate influence of recent observations on expectations, which operates through judgment about the mean of the process. We set up and solve the model in Sections 5.1 and 5.2. We show the basic comparative statics regarding the degree of overreaction in Section 5.3 (and discuss extensions for underreaction). We then assess the model fit in Section 5.4.

5.1 Setup

Environment. Time is discrete and is indexed by $t \in \{0, 1, 2, \dots\}$. There is an agent who forecasts future realizations of an exogenous stochastic process $\{x_t : t \geq 0\}$ at horizon h . The process is AR(1) with mean μ and persistence ρ :

$$x_t = (1 - \rho)\mu + \rho x_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2). \quad (5.1)$$

The agent's payoff at any given time t depends on the accuracy of these forecasts and is given by: $-(F_t x_{t+h} - x_{t+h})^2$, where $F_t x_{t+h}$ is the agent's time t forecast of x 's realization h periods ahead and x_{t+h} is the *ex post* realization of the variable at $t + h$.¹⁷ The model can be extended to general Gaussian ARMA processes and the qualitative conclusions of the model are unchanged, which we show in Online Appendix IA4.

Information Processing. We assume that the agent is uncertain about the long-run mean (μ), but can process information to form an assessment of its value in order to forecast x_{t+h} .

¹⁷Note that x_{t+h} is not fully known at time t and only realized h periods after the forecast is made. Nonetheless, at time t , the agent knows that the payoff is determined by the realization of the process at $t + h$. This is similar to the score function in the experiment. A minor difference is that the score function in the experiment does not have an exact quadratic form to ensure that payments in the experiment are always non-negative (as discussed in Section 3.2). We use this standard quadratic form for simplicity of modeling, so we can derive closed-form solutions.

We model this using information processing where more recent observations are less costly to process than others. The simplest way to obtain such a cost structure is to assume that the agent processes the most recent observation x_t for free, and processes further information at a cost.¹⁸ Formally, after observing x_t , the agent automatically forms the initial prior regarding $\mu \sim N(x_t, \underline{\tau}^{-1})$. The agent then decides whether to process more information to update this prior. Letting S_t denote the set of all processed information (which includes x_t), we assume the cost of processing S_t is increasing and convex in the amount of information processed conditional on x_t :

$$C_t(S_t) \equiv \omega \frac{\exp(\gamma \cdot \mathbb{I}(S_t, \mu|x_t)) - 1}{\gamma},$$

where $\omega \geq 0$ and $\gamma \geq 0$ are the scale and convexity parameters, and $\mathbb{I}(S_t, \mu|x_t)$ is Shannon’s mutual information function, which measures the amount of information used by the agent after observing x_t .¹⁹ We use the term “on top of mind” to refer to the set of information actively used by the agent, S_t . In the extreme case when processing any additional information is costless, our model nests the frictionless rational benchmark. However, with costly information processing, only a subset of data will be on top of mind.

Psychological Foundations. Our model has two key assumptions. First, the agent faces frictions in processing information and use a subset of the data available to them. This assumption is closely related to the psychology literature on working memory, which emphasizes that some information is more actively used than others. This notion has been referred to as heightened activation, increased accessibility, or (constrained) focus of attention (Baddeley and Hitch, 1993; Cowan, 1998, 2017a). Here, we use the term “on top of mind” to refer to the set of information that are at a state of heightened activation (S_t defined above). The psychology literature shows that heightened activation applies even when there is no

¹⁸More generally, one can allow for more observations to be processed for free or all observations to be costly to process. Our model predictions are largely unchanged as long as information processing is relatively cheaper for recent observations than for others. In line with our main assumption, Table A.7 shows that x_t has a disproportionate impact on both the forecast and the deviation from the rational benchmark in our data.

¹⁹This functional form embeds two useful cases. First, it becomes linear in $\mathbb{I}(S_t, \mu|x_t)$ when $\gamma \rightarrow 0$, which is the classic formulation in rational inattention (Sims, 2003). Second, in case of Gaussian beliefs with $\gamma > 1$, the cost is equivalent to choosing the precision of beliefs about μ (see Online Appendix IA2.2 for formal derivations).

explicit recall involved (Spillers, Brewer and Unsworth, 2012).²⁰ Even though a person may see a large number of observations in the environment (e.g., many data points on the screen), not all information is necessarily actively processed. As Barrett, Tugade and Engle (2004) write: “any environment contains an array of stimuli ... but stronger representations will laterally inhibit weaker ones, and the strongest will be expressed in behavior.”²¹

Second, the agent relies on recent information, which is easier to use, and additional information processing, which is subject to a cognitive cost. We include these two components in S_t . This follows the psychology literature’s view about the two ways for information to enter working memory, perhaps best summarized by Hitch, Hu, Allen and Baddeley (2018): “Previous research ... indicates that access to the focus of attention (FoA) can be achieved in either of two ways. The first is automatic and is indexed by the recency effect, the enhanced retention of the final item. The second is strategic and based on instructions to prioritize items differentially, a process that draws on executive capacity and boosts retention of information deemed important.” Furthermore, these two forces in the working memory mechanism correspond to broader themes in psychology research about dual processing: recency (which is part of the information that is automatically activated) and goal-driven information processing (which is slower and requires effort). Thus, our model also relates to a broad class of dual process models in psychology (Barrett, Tugade and Engle, 2004), which is unified by a framework where the individual starts from a default driven by what is immediately accessible (“System 1”), and further adjusts beliefs by effortful processing (“System 2”).²²

Online Appendix IA3 provides a more detailed review of the relevant psychology literature. The mechanism in our model can reflect costly information processing as discussed above, but more generally it can also capture other psychological and institutional factors that limit the usage of past data and generate recency effects.

²⁰The working memory literature emphasizes that heightened activation is not necessarily about recall but rather about focusing on a subset of the available information in processing. See, e.g., Unsworth and Spillers (2010) for a discussion. Cowan (2017a) also explains that the term “working memory” originated from the use in computer science (where it refers to holding information for active processing).

²¹The recent survey paper by Cowan (2017a) explains the concept of working memory as “the ensemble of components of the mind that hold information temporarily in a heightened state of availability for use in ongoing information processing.” Cowan (2017b) conveys the idea through its title, “Working Memory: The Information You Are Now Thinking of.”

²²See Evans (2008) for a summary of the many types of dual process models in the psychology literature and Ilut and Valchev (2023) for an application of dual processing in economics.

Forecasts. Given the estimate of the long-run mean conditional on S_t , $\mathbb{E}[\mu|S_t]$, the agent's forecast of $F_t x_{t+h}$ is:

$$F_t x_{t+h} = \rho^h x_t + (1 - \rho^h) \mathbb{E}[\mu|S_t] = \underbrace{E_t x_{t+h}}_{\text{rational forecast}} + \underbrace{(1 - \rho^h)(\mathbb{E}[\mu|S_t] - \mu)}_{\text{forecast error of long-run mean}}. \quad (5.2)$$

We assume that the agent uses the correct ρ for simplicity. As discussed earlier, modeling frictions in beliefs about the long-run mean μ is the most parsimonious way to capture how overreaction varies with the process persistence and forecast horizon, whereas biases in ρ by itself are not sufficient (as shown by the results across different forecast horizons). Overall, we do not rule out that forecasters may also use an incorrect ρ . Nonetheless, we find that modeling biases about the mean μ is the most concise approach to capture the empirical evidence that overreaction is stronger both when the process is less persistent and when the forecast horizon is longer.

5.2 Model Solution

Given the primitives of the problem at time t , the agent solves:

$$\begin{aligned} \min_{S_t} \mathbb{E} \left[\min_{F_t x_{t+h}} \mathbb{E} \left[(F_t x_{t+h} - x_{t+h})^2 | S_t \right] + C_t(S_t) \right] & \quad (5.3) \\ \text{s.t.} \quad \underbrace{\{x_t\}}_{\text{recent observation}} & \subseteq \underbrace{S_t}_{\text{actively used information}} \subseteq \underbrace{\mathcal{A}_t}_{\text{total information}} \end{aligned}$$

where \mathcal{A}_t is the largest possible set of information that is available for processing given the set of available observations $x^t \equiv \{x_\tau\}_{\tau \leq t}$.²³ In Online Appendix IA2.2, we show that the above problem can be simplified to choosing the optimal precision of the long-run mean estimate:

²³Formally, $\mathcal{A}_t \equiv \{s | \mathbb{I}(s, \mu | x^t) = 0\}$, meaning that no available signal should contain further information about the long-run than what is revealed by the history of available observations.

Lemma 1. *The agent's problem in Equation (5.3) is equivalent to:*

$$\min_{\underline{\tau} \leq \tau \leq \bar{\tau}_t} \left\{ \underbrace{\frac{(1 - \rho^h)^2}{\tau}}_{\text{benefit of precision}} + \omega \underbrace{\frac{\left(\frac{\tau}{\underline{\tau}}\right)^\gamma - 1}{\gamma}}_{\text{cost of precision}} \right\}, \quad (5.4)$$

where $\tau \equiv \text{var}(\mu|S_t)^{-1}$ is the precision of the agent's posterior belief about the long-run mean, and $\bar{\tau}_t$ is the maximum precision obtainable given the full information set \mathcal{A}_t , and $\underline{\tau}$ is the precision of the agent's prior about the long-run mean.

For the remainder of the paper, we assume that the constraint $\tau \leq \bar{\tau}_t$ does not bind, which occurs when \mathcal{A}_t is sufficiently large. As formally shown in Online Appendix IA2.3, it is straightforward to derive the optimal posterior precision of the long-run mean, $\tau^* = \text{var}(\mu|S^t)^{-1}$, as $\tau^* = \underline{\tau} \max \left\{ 1, \left(\frac{(1 - \rho^h)^2}{\omega \underline{\tau}} \right)^{\frac{1}{1+\gamma}} \right\}$, with the agent's forecast error of the long-run mean given by:

$$\mathbb{E}[\mu|S_t] - \mu = \frac{\tau}{\tau^*} (x_t - \mu) + \text{noise}. \quad (5.5)$$

Equation (5.5) shows that the agent's forecast error of the long-run mean is anchored towards x_t , which follows from the assumption that x_t is easier to process than other information. The dependence of the agent's perception of the long-run mean on the most recent observation is the key force that drives the overreaction in forecasts. By applying Equation (5.5) to Equation (5.2), we obtain the following proposition, which presents the behavior of forecasts under the optimal information processing.

Proposition 1. Forecasts display systematic overreaction relative to the rational benchmark, with

$$F_t x_{t+h} = \underbrace{E_t x_{t+h}}_{\text{rational forecast}} + \underbrace{(1 - \rho^h) \min \left\{ 1, \left(\frac{\omega \underline{\tau}}{(1 - \rho^h)^2} \right)^{\frac{1}{1+\gamma}} \right\}}_{\text{overreaction } (\equiv \Delta)} x_t + \underbrace{u_t}_{\text{noise}}. \quad (5.6)$$

Proof. See Online Appendix IA2.3. □

Proposition 1 shows that in presence of costly information processing, forecasts are *more sensitive* to the most recent observation relative to the rational forecasts formed under full

information. We denote this excess sensitivity by the term Δ , which captures the degree of overreaction in forecasts relative to rational forecasts.

5.3 Comparative Statics

We now illustrate the implications of our model for the empirical evidence on overreaction. As Proposition 1 shows, forecasts display overreaction in our model since the agent overweighs the most recent observation x_t when inferring the long-run mean μ .²⁴ Furthermore, as the agent conflates some part of the transitory shock as permanent, we also obtain predictions about how the degree of overreaction varies across different settings. The key intuition is as follows. First, overreaction is stronger when the most recent shock is less predictive of future outcomes. Accordingly, our model naturally generates greater overreaction for less persistent processes and for longer-horizon forecasts. Second, as Equation (5.6) shows, the response of the forecast to x_t in our model also adapts partially to the true process (in line with what we see in the data). In other words, costly information processing allows the agent to moderate the influence of x_t on the forecast. Therefore, the implied persistence is lower when actual ρ is smaller, instead of being a constant, but the adjustment is imperfect.

The following proposition summarizes the combined effect of these two forces and provides comparative statistics of overreaction with respect to the parameters of the model.

Proposition 2. Consider the excess sensitivity of forecasts to x_t measured by $\Delta = \rho_{s,h} - \rho^h$ defined in Equation (5.6):

1. $\Delta \geq 0$ with $\Delta = 0$ if and only if, either $\rho = 1$, or information processing is frictionless ($\omega = 0$) and past information available to the forecaster is infinite.
2. Δ is increasing in $\underline{\tau}$ and ω .
3. Δ is decreasing in ρ^h (and hence decreasing in ρ and increasing in h) if the cost function is weakly convex in τ , which is true if and only if $\gamma \geq 1$.

²⁴This is a fundamental difference between our model and models of sticky information (which may use similar modeling techniques). In sticky information models, agents have full access to past information, but some may not have access to the most recent observation. Accordingly, forecasts can exhibit underreaction (since they rely more on past information rather on the recent observation). In contrast, in our model, the agent is fully aware of the most recent observation but processes past information imperfectly, which results in overreaction (since forecasts rely more on the recent observation than on past information).

4. Denote $\zeta(\rho, h) \equiv \rho_{s,h}/\rho^h$. Then, $\zeta(\rho, h) \geq 1$. Like Δ , $\zeta(\rho, h)$ is decreasing in ρ and increasing in h for all values of ρ and h , if and only if $\gamma \geq 1$.

Proof. See Online Appendix [IA2.4](#). □

Proposition 2 implies that the excess sensitivity of forecasts to the most recent observation (as measured by Δ or $\zeta(\rho, h)$), is decreasing in the persistence ρ and increasing in the horizon h , as long as the cost of information processing is weakly convex in the precision τ . Moreover, this insight carries over to the gap between the implied persistence ρ_h^s and the true persistence ρ , especially for less persistent processes.²⁵ We derive this formally in the following corollary and show the numerical results in Figure VI.

Corollary 1. *Consider the measure of the implied persistence per period relative to the actual persistence $\Delta\rho_h \equiv \rho_h^s - \rho$, and also assume $\gamma \geq 1$. Then, $\Delta\rho_h \geq 0$ and decreasing in ρ . Moreover, $\Delta\rho_h$ is increasing in h for h sufficiently large: holding fixed ρ , there exists an h^* such that $\Delta\rho_h$ is increasing in h for $h \geq h^*$.*

Proof. See Online Appendix [IA2.5](#). □

In sum, Proposition 2, along with Corollary 1, delivers two main results of our model. The first result is overreaction, a prediction that is consistent with the evidence presented in Section 4: the gap between the implied persistence and the actual persistence, Δ and $\Delta\rho_h$, is positive (or equivalently, $\zeta(\rho, h)$ is greater than 1). The second result is that overreaction (as measured by Δ , $\Delta\rho_h$, or $\zeta(\rho, h)$) is stronger for less persistent processes and sufficiently long horizons, as we observe in the data (as long as the cost of information utilization is convex in the precision of the agent's forecast).²⁶

²⁵More precisely, we show in Online Appendix [IA2.5](#) that $\Delta\rho_h$ is monotonically decreasing in ρ^h for $\rho^h \leq \lambda$, a positive constant independent of ρ and h which depends on γ and $\omega\underline{\tau}$. In practice, as shown in Figure VI, we find $\Delta\rho_h$ is increasing in the range of persistence and horizon covered in our experiment for our calibrated parameters.

²⁶Another approach for modeling overreaction is the one in [da Silveira, Sung and Woodford \(2020\)](#) who introduce a dynamic framework where memory is costly and agents optimally choose their memory structure over time. In their model, agents decide what they want to remember in the future before an observation is realized. In our model, the recent observation is the starting point and agents decide to use past information after an observation has been realized. While both our model and the model in [da Silveira, Sung and Woodford \(2020\)](#) deliver overreaction in posterior beliefs, the predictions for prior beliefs are different: in our model, the priors are anchored to the present, as the utilization of past data happens after the most recent observation is realized; in [da Silveira, Sung and Woodford \(2020\)](#), in contrast, the priors depend on the memory state which is formed optimally based on the past observations before the most recent observation is realized.

Finally, the baseline version of the model focuses on overreaction in light of our empirical evidence presented in Section 4. We provide an extension in Online Appendix IA5 that shows how the model can allow for underreaction as well. In particular, if the signals are noisy (Woodford, 2003) or if updating is infrequent (Mankiw and Reis, 2002), then there can be an additional force that pushes in the direction of underreaction. In this case, the model shows that overreaction will still be relatively more pronounced when the process is less persistent. In our experiment, the signal is clear and infrequent updating is unlikely, so overreaction dominates.

5.4 Model Fit

In addition to the qualitative comparative statics presented above, we now estimate the model using our forecast data to further evaluate its performance. We present results on model fit for the case where the cost of information utilization is quadratic ($\gamma = 2$). We also present robustness checks in Section 6.2 where we jointly estimate γ with the other parameters of the model, which produce similar results. As before, the model is estimated by minimizing the mean-squared error (MSE) between the one-period forecast predicted by the model for a given parameter (using the realizations of x_t in the data) and the one-period forecast observed in the data.

First, we show that our model matches the relationship between the implied persistence and the actual persistence found in the data. Figure VI, Panel A, shows the results for the baseline horizon $h = 1$. The solid line represents the implied persistence ρ_1^s in the data, and the red solid circles represent ρ_1^s predicted by our model. We see that the implied persistence ρ_1^s predicted by our model is very similar to that in the data. Note that there is nothing mechanical in this very good fit. The models investigated in Figure IV were fitted the same way as our model, and do not fit the empirical relation between ρ_1^s and ρ .

We also examine model fit for longer-horizon forecasts in Figure VI, Panels B to D. We present the implied persistence (per period) ρ_h^s for horizons $h = 2, 5, 10$. Importantly, we fit our model using one-period-ahead forecasts, so its performance for other forecast horizons is not targeted. We see that our model performs well for all forecast horizons.

Table A.8 further evaluates the model fit by calculating the MSE between ρ_h^s in the model

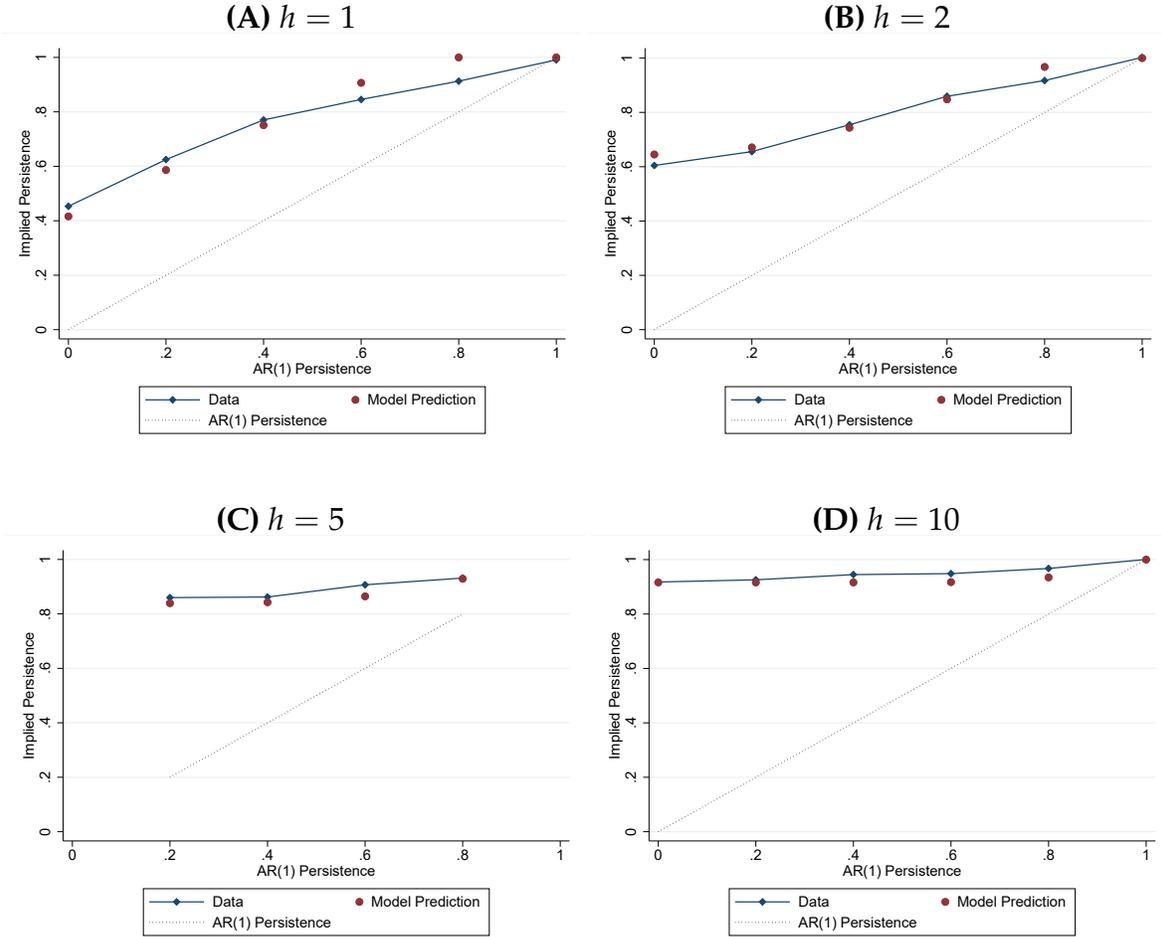


FIGURE VI

Model Fit: Implied Persistence

This figure shows the forecast-implied persistence ρ_h^s as a function of the actual persistence ρ . The implied persistence ρ_h^s is obtained by regressing $F_t x_{t+h}$ on x_t and taking the regression coefficient to the $(1/h)$ th power. The blue line represents the results in the forecast data. The solid red dot represents ρ_1^s from our model. Results for $h = 1$ and $h = 2$ use forecasts in Experiment 1. Results for $h = 5$ and $h = 10$ use forecasts in Experiment 2 and Experiment 4, respectively. We fit the model parameter using data for $h = 1$ and use the same parameter to generate model predictions for other forecast horizons.

and ρ_h^s in the data, as well as the MSE between $F_t x_{t+h}$ in the model and in the data. We calculate the MSE for our model and the models in Section 4.2. This MSE calculation also confirms that our model performs well.

6 Additional Tests and Robustness

In Section 6.1, we present further experimental results to shed light on the mechanisms in the model. In Section 6.2, we show the robustness of our model formulations to different functional forms, and we discuss several modeling assumptions.

6.1 Further Experimental Results

We design additional experiments to investigate the mechanism in our model.

Bias about the long-run mean The key component of our model is that biases in the forecasts are linked to biases in the long-run mean assessment, which loads too much on the recent observation relative to the frictionless benchmark as shown in Equation (5.5). We test this mechanism in an additional experiment, where we elicit participants' assessment of the long-run mean by asking them to enter a response in every round for the following question: "What do you think is the long-run average of the process?" In the same experiment, participants also make long-horizon forecasts about x_{t+10} in each round, so we can compare the consistency between the x_{t+10} forecast and the long-run mean assessment. A screenshot of this experimental interface is presented in Panel A of Figure A.9. We randomly assign each participant to a given value of AR(1) persistence ρ , and we focus on conditions with $\rho \in \{0, 0.2, 0.4, 0.6, 0.8\}$ (as the long-run average is not well-defined for $\rho = 1$). The data is collected in Experiment 5 (conditions E11 to E16 in Table A.2).

We document two main findings that align with our model. First, the long-run mean assessment ($F_t\mu$) helps explain the x_{t+10} forecast. For example, if we use the actual value of ρ , then we can impute the long-run mean implied by the x_{t+10} forecast: $\tilde{\mu}_t = \frac{F_t x_{t+10} - \rho^{10} x_t}{1 - \rho^{10}}$. The difference between the reported long-run mean belief ($F_t\mu$) and the forecast-implied long-run mean ($\tilde{\mu}_t$) has median (mean) 0 (-1), with an interquartile range of -14 to 12.

Columns (1) and (2) of Table I show that the forecast-implied long-run mean ($\tilde{\mu}_t$) is strongly correlated with the long-run mean assessment ($F_t\mu$). In addition, for conditions with $\rho < 1$, we have $\rho^{10} \approx 0$ so the x_{t+10} forecast itself should be reasonably close to the long-run mean assessment ($F_t\mu$).²⁷ In the data, the raw difference between the long-run

²⁷Even $0.8^{10} \approx 0.11$, where 0.8 is the highest persistence in our experimental conditions except $\rho = 1$.

mean belief ($F_t\mu$) and the long-run forecast (F_tx_{t+10}) has median (mean) 0 (3) and interquartile range -15 to 18. Columns (3) and (4) of Table I show that the x_{t+10} forecast is strongly correlated with the long-run mean assessment.

Second, the long-run mean assessment ($F_t\mu$) is significantly influenced by the most recent observation x_t (the forecast of x_{t+10} overreacts to x_t as well, with magnitude similar to the baseline results in Panel C of Figure III).²⁸ In columns (5) and (6) of Table I, we show the loading of the long-run mean assessment ($F_t\mu$) on x_t . We observe a clear positive relationship consistent with the predictions of our model. Taken together, these results are difficult to explain with alternative models. As shown in Section 4.2, commonly used models do not generate enough overreaction for the long-term forecasts, especially when the process is transitory.²⁹ Our model generates overreaction for both the long-term forecast (such as F_tx_{t+10}) and the long-run mean assessment.

A related question is whether we can eliminate forecasting biases if we tell participants the long-run mean of the process. To investigate this question, we performed an experiment where we provided the following information on the instruction page and the prediction page: “The mean of the process is zero. In other words, the long-run mean is zero.” A screenshot of the experimental interface is presented in Panel B of Figure A.9. We randomly assign each participant to a given value of AR(1) persistence ρ , and we focus on conditions with $\rho \in \{0, 0.2, 0.4, 0.6, 0.8\}$ as before. The data is collected in Experiment 5 (conditions E21 to E36 in Table A.2). Table A.9 shows the comparison of this set of treatment conditions with the baseline conditions (participants are randomly assigned to one type of condition). We observe less overreaction in these treatment conditions compared to the baseline, but the biases are not eliminated. A limitation of this test is that the abstract information provided to participants about the long-run mean may not be internalized by them, in which case it could be difficult to make a difference to their forecasting behavior. Future research can explore whether educating participants about the long-run mean in more detail could further eliminate biases.

²⁸These different designs also verify that overreaction in the x_{t+10} forecast is not affected by whether the short-term x_{t+1} forecast is asked.

²⁹These models do not always have a well-defined counterpart to the long-run mean forecast. It is easiest to think that the long-run mean forecast implied by these models is approximated by the forecast of x_{t+h} where h is a relatively large number. In this case, we know that these models do not generate enough overreaction, especially when the process is transitory, as we have seen from the x_{t+10} forecasts.

TABLE I
LONG-RUN MEAN BELIEF

	$\tilde{\mu}_t$		$F_t x_{t+10}$		Long-Run Mean Belief	
	(1)	(2)	(3)	(4)	(5)	(6)
Long-Run Mean Belief	0.604*** (0.085)	0.639*** (0.075)	0.582*** (0.082)	0.629*** (0.073)		
x_t					0.160*** (0.034)	0.180*** (0.029)
Participant FE	N	Y	N	Y	N	Y
Observations	6,095	6,095	6,095	6,095	6,095	6,095
R ²	0.10	0.59	0.10	0.57	0.02	0.49

Notes. In columns (1) and (2), we compare the implied long-run mean in the forecast of x_{t+10} ($\tilde{\mu}_t$) with the long-run mean assessment reported by participants. In columns (3) and (4), we compare the forecast of x_{t+10} ($F_t x_{t+10}$) with the long-run mean assessment reported by participants. In columns (5) and (6), we examine the loading of the long-run mean assessment on the recent observation. The data is collected in Experiment 5. Each participant is randomly assigned to a given ρ . Columns (2), (4), and (6) include participant fixed effects to control for average optimism. Standard errors clustered by participant are presented in parentheses. *** indicates a 1% level of significance.

On top of mind A further component of our model is that forecasts overreact because recent observations are on top of mind and more easily available. Accordingly, forecasters rely excessively on recent observations in their judgment of the long-run mean. To test this mechanism, we implement two conditions that aim to reduce the extent to which the last observation is on top of mind, which should in turn reduce the degree of overreaction. In the first condition (“click x_{t-10} ”), we require participants to click on x_{t-10} before making their forecasts in each round. In the second condition (“red line”), where we draw a red line at zero (the actual mean of the process) on the graphical interface for forecasting. A screenshot of the interface for these two conditions is presented in Panels C and D of Figure A.9. The two treatment conditions seek to divert the focus away from the most recent observation, which can reduce its impact on the assessment of the long-run mean.³⁰ We also include the baseline treatment condition (same design as Experiment 1) for comparison. We focus on conditions with $\rho \in \{0, 0.2, 0.4, 0.6, 0.8\}$ (our model predicts that forecasts do not overreact when $\rho = 1$, which is consistent with the data). Each participant is randomly assigned to a

³⁰Formally, both treatments can be modeled as bringing in other signals about the long-run mean to top of mind, in addition to the existing default belief $\mu \sim N(x_t, \underline{\tau}^{-1})$. We then obtain the prediction that the implied persistence in the new treatment conditions should be smaller. See Online Appendix IA6 for details.

TABLE II
CHANGING WHAT IS ON TOP OF MIND

	$x_{t+1} - F_t x_{t+1}$		$\rho x_t - F_t x_{t+1}$	
	(1)	(2)	(3)	(4)
x_t	-0.267*** (0.033)	-0.334*** (0.029)	-0.272*** (0.031)	-0.292*** (0.028)
× Click on x_{t-10}	0.078 (0.051)	0.109** (0.050)	0.091* (0.049)	0.124** (0.048)
× Red line at $x = 0$	0.099** (0.046)	0.134*** (0.042)	0.101** (0.045)	0.127*** (0.040)
Participant FE	N	Y	N	Y
Observations	18,560	18,560	18,560	18,560
R ²	0.02	0.22	0.04	0.32

Notes. In this table, we regress different definitions of the forecast error (realization minus forecast) on the last realization, interacted with two indicator variables that equal one when the participant is allocated to the new treatment conditions. One of these conditions requires participants to click on the point corresponding to x_{t-10} in each round before entering new forecasts. The other one features a red line at $x = 0$. We also include the baseline condition (same design as Experiment 1) for comparison. The data is collected in Experiment 4. Each participant is randomly assigned to a given ρ and a given condition. Columns (2) and (4) include participant fixed effects to control for average optimism. In all regressions, we exclude conditions for which $\rho = 1$, since in this case we know that forecasts do not have significant biases (the implied persistence is close to one). Standard errors clustered by participant are presented in parentheses. *** indicates a 1% level of significance.

given ρ and a given treatment condition. The data is collected in Experiment 4 (conditions D11 to D36 in Table A.2).

We present the results in Table II using the following regression:

$$x_{it+1} - F_t x_{it+1} = \alpha x_{it} + \beta T_i^{\text{Click } x_{t-10}} \times x_{it} + \gamma T_i^{\text{Red Line}} \times x_{it} + a_i + \epsilon_{it}, \quad (6.1)$$

where $T_i^{\text{Click } x_{t-10}}$ and $T_i^{\text{Red Line}}$ are indicator variables that equal one if individual i is assigned to one of the new treatment conditions, and ϵ_{it} is clustered by forecaster. Since there is strong overreaction in the experiment, we expect $\alpha < 0$. But we expect overreaction to be less pronounced in both conditions where the last observation is less on top of mind, so that $\beta > 0$ and $\gamma > 0$. In the data, we find that both treatments reduce overreaction, in line with the prediction of our model.

6.2 Robustness of Model Formulations

We discuss several main assumptions in our baseline model in Section 5.

A. Convexity and General Functional Form

Our benchmark calibration assumes that the cost of information utilization is quadratic ($\gamma = 2$) in the relative precision $\frac{\tau}{\underline{\tau}}$. Here, we examine two alternative ways for calibrating γ and show the robustness of the results. First, we fit our model assuming the cost is linear in the mutual information ($\gamma \mapsto 0$), which is a standard approach in the rational inattention literature (e.g. Sims, 2003). Second, we fully optimize over the convexity parameter γ using a grid-search method. Figure A.10 shows the results. The linear approach does a reasonable job fitting the implied persistence, but overshoots slightly for processes with higher persistence and undershoots slightly for processes with lower persistence. The general γ approach produces very good fit (the optimal value of γ roughly equals 10). Overall, we find that the model has good performance and is not very sensitive to the exact value used for γ .

B. Assumptions on $\underline{\tau}$

Our main model defines $\underline{\tau}$ as the baseline precision the agent has regarding the long-run mean after seeing the most recent observation. For simplicity, we have assumed that $\underline{\tau}$ is fixed across all experiments and across different persistence levels ρ . In the following, we also consider an alternative approach, where we endogenize $\underline{\tau}$. One natural candidate for $\underline{\tau}$ is the inverse of the variance of the stationary distribution for the AR(1) process:

$$\underline{\tau}^{alt} = \frac{1 - \rho^2}{\sigma_\epsilon^2}. \quad (6.2)$$

This choice can have a Bayesian interpretation as the posterior variance given x_t , for a Bayesian with an improper uniform prior (or a sequence of priors that become increasingly dispersed). In particular, $\underline{\tau}^{alt}$ is decreasing in ρ : the agent is ex ante more uncertain about the long-run mean when the process is unconditionally more volatile. Figure A.11 shows the fit of the alternative specification and confirms the model performs well in this case too.

C. Incentives and External Validity

A possible question is whether one can test the effect of changing forecasters' incentives, or the trade-off between the cost of information processing and the benefit of obtaining accurate beliefs. We performed the following experiments to vary incentives within our budget constraints. The first set of treatment conditions start with 40 rounds of predictions with smaller monetary incentives, where the average bonus is on the scale of \$0.50. (bonus equal to total score divided by 4,000). Then, there is a second round of 40 predictions with larger monetary incentives, where the average bonus is on the scale of \$5.00 (total score divided by 400). The second set of treatment conditions reverse the order, with larger monetary incentives before smaller monetary incentives (to make sure any results are not driven by the effects of order). The data is collected in Experiment 5 (conditions E41 to E46 in Table A.2).

In Panel A of Table A.10, we run the following regression on these two treatment conditions:

$$x_{t+1} - F_t x_{t+1} = \alpha x_t + \beta x_t \times 1_{\text{High Incentive}} + \text{constant},$$

where the dummy $1_{\text{High Incentive}}$ is equal to one for the forecasting rounds with larger incentives. We exclude $\rho = 1$ since overreaction does not occur here in the data and in the model. In these regressions, we have $\alpha < 0$ as usual: participants overreact, on average, to the most recent observation. Meanwhile, we do not find that β is significant. As an additional check, we regress participants' forecasting scores on the incentive regime in Panel B of Table A.10. For each participant, we calculate the total forecasting score during the 40 rounds with larger incentives and the 40 rounds with smaller incentives. We regress the score on a dummy for high incentive rounds. Columns (1) and (2) use the score; columns (3) and (4) use the log score. Scores are slightly higher in the high incentive rounds, but the difference is small and insignificant.

Taken together, consistent with DellaVigna and Pope (2017), we find that variations in experimental incentives have limited effects on forecasting behavior or forecasting scores. As DellaVigna and Pope (2017) discuss, participants' behavior in experiments is not entirely driven by monetary incentives. The variations of experimental payments are inevitably limited by our budgets, and variations in financial rewards within the range of affordability do not seem to have a major impact on participants' performance.

Another question is whether incentives for accuracy in practice could be so large that decision makers will overcome all costs of information utilization. A large literature documents biases in high-stake settings (Malmendier and Tate, 2005; Pope and Schweitzer, 2011; Ben-David, Graham and Harvey, 2013; Greenwood and Hanson, 2015; Bordalo, Gennaioli, La Porta and Shleifer, 2019), which indicate that frictions may not be fully eradicated in these situations. Furthermore, many decisions are made under time constraints or with a fair bit of human discretion, in which case the frictions represented by our model—namely, certain information is particularly on top of mind—are likely to be present. Overall, our findings align with suggestive evidence from field data as discussed earlier; across many settings, we observe consistent patterns of stronger overreaction when the process persistence is lower and the forecast horizon is longer.

7 Conclusion

Recent research using survey data from different sources points to varying degrees of biases in expectations. A key question is how to unify these different findings. To have a better understanding of how biases vary with the setting, we conduct a large-scale randomized experiment where participants forecast stable random processes. The experiment allows us to control the DGP and the relevant information sets. This is not feasible in survey data, which can give rise to major complications in interpreting the results.

We find that forecasts display significant overreaction: they respond too much to recent observations. Furthermore, overreaction is particularly pronounced for less persistent processes and longer forecast horizons. We also find that standard specifications of commonly-used expectations models, estimated in our data, do not square with the variation in overreaction. Some predict too much overreaction when the process is transitory, while others predict too little.

We propose a framework to capture biases in expectations formation, where forecasters form estimates of the long-run mean of the process using a mix of the recent observation and past data. They balance these two sources of information depending on the setting, under the constraint that the utilization of past information is costly. As a result, forecasts adapt partially to the setting, but recent observations have a disproportionate influence, resulting

in overreaction. Over-adjusting the estimates of the long-run mean in response to recent observations also implies that overreaction is more pronounced when the process is more transitory and the forecast horizon is longer. We estimate the model in our data and find that it closely matches how overreaction varies with process persistence. The model, when estimated on short-term forecasts, also predicts biases in long-term forecasts that closely match what we observe in the data.

Finally, our baseline model focuses on overreaction given the empirical evidence in our experiment. Nonetheless, the model can also be extended to allow for underreaction by introducing noisy signals, which could be a reason for underreaction observed in some survey data (Coibion and Gorodnichenko, 2012; Bouchaud et al., 2019). In this setting, the model maintains the prediction that the degree of overreaction should be relatively stronger when the process is less persistent. Taken together, we hope that the evidence and theory in this paper contribute to a systematic understanding of the findings on expectation biases.

References

- Abarbanell, Jeffrey and Victor Bernard**, "Tests of analysts' overreaction/underreaction to earnings information as an explanation for anomalous stock price behavior," *Journal of Finance*, 1992.
- Amromin, Gene and Steven A Sharpe**, "From the horse's mouth: Economic conditions and investor expectations of risk and return," *Management Science*, 2013, 60 (4), 845–866.
- Angeletos, George-Marios, Zhen Huo, and Karthik A Sastry**, "Imperfect macroeconomic expectations: Evidence and theory," *NBER Macroeconomics Annual*, 2021, 35 (1), 1–86.
- Assenza, Tiziana, Te Bao, Cars Hommes, and Domenico Massaro**, "Experiments on expectations in macroeconomics and finance," *Research in Experimental Economics*, 2014, 17, 11–70.
- Baddeley, Alan D and Graham Hitch**, "The recency effect: Implicit learning with explicit retrieval?," *Memory & Cognition*, 1993, 21 (2), 146–155.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny**, "A model of investor sentiment," *Journal of Financial Economics*, 1998, 49 (3), 307–343.
- , **Robin Greenwood, Lawrence Jin, and Andrei Shleifer**, "X-CAPM: An extrapolative capital asset pricing model," *Journal of Financial Economics*, 2015, 115, 1–24.
- Barrero, Jose Maria**, "The micro and macro of managerial beliefs," *Journal of Financial Economics*, 2021.
- Barrett, Lisa Feldman, Michele M Tugade, and Randall W Engle**, "Individual differences in working memory capacity and dual-process theories of the mind," *Psychological Bulletin*, 2004, 130 (4), 553.
- Ben-David, Itzhak, John R Graham, and Campbell R Harvey**, "Managerial miscalibration," *Quarterly Journal of Economics*, 2013, 128 (4), 1547–1584.
- Benjamin, Daniel J**, "Errors in probabilistic reasoning and judgment biases," *Handbook of Behavioral Economics: Applications and Foundations*, 2019, 2, 69–186.
- Beshears, John, James J Choi, Andreas Fuster, David Laibson, and Brigitte C Madrian**, "What goes up must come down? Experimental evidence on intuitive forecasting," *American Economic Review*, 2013, 103 (3), 570–574.
- Bloomfield, Robert and Jeffrey Hales**, "Predicting the next step of a random walk: Experimental evidence of regime-shifting beliefs," *Journal of Financial Economics*, 2002, 65 (3), 397–414.
- Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon, and Andrei Shleifer**, "Memory and probability," *Quarterly Journal of Economics*, 2023, 138 (1), 265–311.
- , **Katherine Coffman, Nicola Gennaioli, Frederik Schwerter, and Andrei Shleifer**, "Memory and representativeness.," *Psychological Review*, 2020, 128 (1), 71.

- , **Nicola Gennaioli**, and **Andrei Shleifer**, “Diagnostic expectations and credit cycles,” *Journal of Finance*, 2018, 73 (1), 199–227.
- , – , and – , “Memory, attention, and choice,” *Quarterly Journal of Economics*, 2020, 135 (3), 1399–1442.
- , – , **Rafael La Porta**, and **Andrei Shleifer**, “Diagnostic expectations and stock returns,” *Journal of Finance*, 2019, 74 (6), 2839–2874.
- , – , **Yueran Ma**, and **Andrei Shleifer**, “Overreaction in macroeconomic expectations,” *American Economic Review*, 2020, 110 (9), 2748–82.
- Bouchaud, Jean-Philippe**, **Philipp Krueger**, **Augustin Landier**, and **David Thesmar**, “Sticky expectations and the profitability anomaly,” *Journal of Finance*, 2019, 74 (2), 639–674.
- Cagan, Phillip**, “The monetary dynamics of hyperinflation,” *Studies in the Quantity Theory of Money*, 1956.
- Casler, Krista**, **Lydia Bickel**, and **Elizabeth Hackett**, “Separate but equal? A comparison of participants and data gathered via Amazon’s MTurk, social media, and face-to-face behavioral testing,” *Computers in Human Behavior*, 2013, 29 (6), 2156–2160.
- Cavallo, Alberto**, **Guillermo Cruces**, and **Ricardo Perez-Truglia**, “Inflation expectations, learning, and supermarket prices: Evidence from survey experiments,” *American Economic Journal: Macroeconomics*, 2017, 9 (3), 1–35.
- Coibion, Olivier** and **Yuriy Gorodnichenko**, “What can survey forecasts tell us about information rigidities?,” *Journal of Political Economy*, 2012, 120, 116–159.
- and – , “Information rigidity and the expectations formation process: A simple framework and new facts,” *American Economic Review*, 2015, 105 (8), 2644–78.
- Cowan, Nelson**, *Attention and memory: An integrated framework*, Oxford University Press, 1998.
- , “The many faces of working memory and short-term storage,” *Psychonomic Bulletin & Review*, 2017, 24 (4), 1158–1170.
- , “Working memory: The information you are now thinking of,” in John H. Byrne, ed., *Learning and Memory: A Comprehensive Reference (Second Edition)*, second edition ed., Oxford: Academic Press, 2017, pp. 147–161.
- da Silveira, Rava Azeredo** and **Michael Woodford**, “Noisy memory and overreaction to news,” in “AEA Papers and Proceedings,” Vol. 109 2019, pp. 557–61.
- , **Yeji Sung**, and **Michael Woodford**, “Optimally imprecise memory and biased forecasts,” NBER Working Paper No. 28075 2020.
- D’Acunto, Francesco**, “Identity and Choice Under Risk,” SSRN Working Paper 2020.

- **and Michael Weber**, “Memory and beliefs: Evidence from the field,” Working Paper 2020.
- d’Arienzo, Daniele**, “Maturity Increasing Overreaction and Bond Market Puzzles,” SSRN Working Paper 2020.
- De Bondt, Werner and Richard H Thaler**, “Do security analysts overreact?,” *American Economic Review*, 1990, pp. 52–57.
- DellaVigna, Stefano and Devin Pope**, “What motivates effort? Evidence and expert forecasts,” *Review of Economic Studies*, 2017, 85 (2), 1029–1069.
- **and** – , “Predicting experimental results: Who knows what?,” *Journal of Political Economy*, 2018, 126 (6), 2410–2456.
- Dwyer, Gerald P, Arlington W Williams, Raymond C Battalio, and Timothy I Mason**, “Tests of rational expectations in a stark setting,” *Economic Journal*, 1993, 103 (418), 586–601.
- Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann**, “Associative memory and belief formation,” NBER Working Paper No. 26664 2020.
- Evans, George and Seppo Honkapohja**, *Learning and Expectations in Macroeconomics*, Princeton University Press, 2001.
- Evans, Jonathan BT**, “Dual-processing accounts of reasoning, judgment, and social cognition,” *Annu. Rev. Psychol.*, 2008, 59, 255–278.
- Fan, Tony, Yucheng Liang, and Cameron Peng**, “The Inference-Forecast Gap in Belief Updating,” SSRN Working Paper 2021.
- Frydman, Cary and Gideon Nave**, “Extrapolative beliefs in perceptual and economic decisions: Evidence of a common mechanism,” *Management Science*, 2016, 63 (7), 2340–2352.
- Fuster, Andreas, David Laibson, and Brock Mendel**, “Natural expectations and macroeconomic fluctuations,” *Journal of Economic Perspectives*, 2010, 24 (4), 67–84.
- Gabaix, Xavier**, “Behavioral inattention,” *Handbook of Behavioral Economics*, 2018.
- Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer**, “Expectations and Investment,” *NBER Macroeconomics Annual*, 2016, 30 (1), 379–431.
- Giglio, Stefano and Bryan Kelly**, “Excess volatility: Beyond discount rates,” *Quarterly Journal of Economics*, 2018, 133 (1), 71–127.
- Greenwood, Robin and Andrei Shleifer**, “Expectations of returns and expected returns,” *Review of Financial Studies*, 2014, 27 (3), 714–746.
- **and Samuel G Hanson**, “Waves in ship prices and investment,” *Quarterly Journal of Economics*, 2015, 130 (1), 55–109.
- Hartzmark, Samuel M, Samuel Hirshman, and Alex Imas**, “Ownership, learning, and beliefs,” *Quarterly Journal of Economics*, 2021, 136 (3), 1665–1717.

- Hey, John D**, “Expectations formation: Rational or adaptive or ...?,” *Journal of Economic Behavior & Organization*, 1994, 25 (3), 329–349.
- Hirshleifer, David, Jun Li, and Jianfeng Yu**, “Asset pricing in production economies with extrapolative expectations,” *Journal of Monetary Economics*, 2015, 76, 87–106.
- Hitch, Graham J, Yanmei Hu, Richard J Allen, and Alan D Baddeley**, “Competition for the focus of attention in visual working memory: Perceptual recency versus executive control,” *Annals of the New York Academy of Sciences*, 2018, 1424 (1), 64–75.
- Ilut, Cosmin and Rosen Valchev**, “Economic agents as imperfect problem solvers,” *The Quarterly Journal of Economics*, 2023, 138 (1), 316–362.
- Kahana, Michael Jacob**, *Foundations of Human Memory*, Oxford University Press, 2012.
- Kahneman, Daniel and Amos Tversky**, “Subjective probability: A judgment of representativeness,” *Cognitive Psychology*, 1972, 3 (3), 430–454.
- Kohlhas, Alexandre and Ansgar Walther**, “Asymmetric attention,” *American Economic Review*, September 2021, 111 (9), 2879–2925.
- Kuziemko, Ilyana, Michael I Norton, Emmanuel Saez, and Stefanie Stantcheva**, “How elastic are preferences for redistribution? Evidence from randomized survey experiments,” *American Economic Review*, 2015, 105 (4), 1478–1508.
- Lian, Chen, Yueran Ma, and Carmen Wang**, “Low interest rates and risk-taking: Evidence from individual investment decisions,” *Review of Financial Studies*, 09 2018, 32 (6), 2107–2148.
- Ma, Yueran, Tiziano Ropele, David Sraer, and David Thesmar**, “A quantitative analysis of distortions in managerial forecasts,” NBER Working Paper No. 26830 2020.
- Malmendier, Ulrike and Geoffrey Tate**, “CEO overconfidence and corporate investment,” *Journal of Finance*, 2005, 60 (6), 2661–2700.
- **and Stefan Nagel**, “Learning from inflation experiences,” *Quarterly Journal of Economics*, 2016, 131 (1), 53–87.
- Mankiw, Gregory and Ricardo Reis**, “Sticky information versus sticky prices: A proposal to replace the New Keynesian Philips Curve,” *Quarterly Journal of Economics*, 2002, 117 (4), 1295–1328.
- Massey, Cade and George Wu**, “Detecting regime shifts: The causes of under- and overreaction,” *Management Science*, 2005, 51 (6), 932–947.
- Metzler, Lloyd A**, “The nature and stability of inventory cycles,” *Review of Economics and Statistics*, 1941, 23 (3), 113–129.
- Nagel, Stefan and Zhengyang Xu**, “Asset pricing With fading memory,” *The Review of Financial Studies*, 2022, 35 (5), 2190–2245.

- Neligh, Nathaniel**, “Rational memory with decay,” SSRN Working Paper 2022.
- Nerlove, Marc**, “Adaptive expectations and cobweb phenomena,” *Quarterly Journal of Economics*, 1958, 72 (2), 227–240.
- Pope, Devin G and Maurice E Schweitzer**, “Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes,” *American Economic Review*, 2011, 101 (1), 129–57.
- Rabin, Matthew**, “Inference by believers in the law of small numbers,” *Quarterly Journal of Economics*, 2002, 117 (3), 775–816.
- **and Dimitri Vayanos**, “The gambler’s and hot-hand fallacies: Theory and applications,” *Review of Economic Studies*, 2010, 77 (2), 730–778.
- Reimers, Stian and Nigel Harvey**, “Sensitivity to autocorrelation in judgmental time series forecasting,” *International Journal of Forecasting*, 2011, 27 (4), 1196–1214.
- Sims, Christopher A**, “Implications of rational inattention,” *Journal of Monetary Economics*, 2003, 50 (3), 665–690.
- Spillers, Gregory, Gene Brewer, and Nash Unsworth**, “Working memory and information processing,” *Encyclopedia of the Sciences of Learning*, New York: Springer, 2012, pp. 3474–3476.
- Unsworth, Nash and Gregory J Spillers**, “Working memory capacity: Attention control, secondary memory, or both? A direct test of the dual-component model,” *Journal of Memory and Language*, 2010, 62 (4), 392–406.
- Wachter, Jessica A and Michael Jacob Kahana**, “A retrieved-context theory of financial decisions,” NBER Working Paper No. 26200 2019.
- Wang, Chen**, “Under- and over-reaction in yield curve expectations,” SSRN Working Paper 2021.
- Woodford, Michael**, “Imperfect common knowledge and the effects of monetary policy,” *Knowledge, Information, and Expectations in Modern Macroeconomics*, 2003.