Understanding Rationality and Disagreement in House Price Expectations

Zigang Li* Stijn Van Nieuwerburgh† Wang Renxuan‡

July 25, 2023

Abstract

Professional house price forecast data are consistent with a rational model where agents must learn about the parameters of the house price growth process and the underlying state of the housing market. Slow learning about the long-run mean can generate forecast bias, a response of forecasts to lagged realizations, sluggish response of forecasts to contemporaneous realizations, and over-reaction to forecast revisions. Introducing behavioral biases, either over-confidence or diagnostic expectations, helps the model further improve its predictions for short-horizon over-reaction and dispersion. Using panel data for a cross-section of forecasters and a term structure of forecasts are important for generating these results.

*Rotman School of Management, University of Toronto, 105 St George St, Toronto, ON M5S 3E6; Email: zigang.li@rotman.utoronto.ca.

†Columbia University Graduate School of Business, NBER, and CEPR, 665 West 130th Street, New York, NY 10027; Email: svnieuwe@gsb.columbia.edu; Web: https://www0.gsb.columbia.edu/faculty/svannieuwerburgh/.

‡CEIBS, 699 Hongfeng Rd, Pudong, Shanghai, China, 201203; Email: rxwang@ceibs.edu.

The authors thank Pulsenomics for making their micro survey data available for academic research. One of the authors is a participant in the Pulsenomics survey. He has not received any compensation for this participation nor has Pulsenomics had any say in the content of this article. The authors thank Alina Arifeva, Hongye Guo (discussant), Zhiguo He, William Larson, Haoyang Liu (discussant), Liyan Yang, Anthony Yezer, and conference participants at the AREUEA National Meeting and the Sixth CEIBS Finance and Accounting Symposium.

1 Introduction

A recent and active literature studies how economic agents form expectations and how those expectations affect outcomes. The housing market is an important laboratory for this inquiry given the prominence of housing in households' portfolios and in the broader functioning of the macroeconomy, and given the role that expectations may play in driving aggregate house price fluctuations.[1]

This research agenda has been propelled by the increasing availability of survey data. In the realm of real estate, survey data on house price forecasts have been available for a while (Case and Shiller, 1988), but their scope continues to improve as discussed in the survey paper by Kuchler, Piazzesi and Stroebel (2023). One finding from surveying households is that past house price growth is an important driver of forecasts of future house price growth. While time series variation in expectations is naturally smaller than time series variation in realizations, house price expectations are nevertheless substantially more optimistic following recent price increases. Households seem to extrapolate from past realizations.[2] Another stylized fact is that the cross-sectional dispersion in forecasts of house price growth exceeds the time-series variation in the average forecast. There has been only limited success in empirically accounting for this heterogeneity.

The survey evidence naturally raises the question of how to best understand the expectation formation process in the housing market. The debate on whether extrapolation represents an optimal use of information, for example, is complicated by the fact that, unlike stock and bond returns, house price changes display substantial serial correlation. In an experimental setting, Armona, Fuster and Zafar (2019) find that subjects revise both short-term and long-term expectations when they receive information about past realized house price changes. They argue that, while respondents extrapolate from past information on average, they do so too little at short horizons (one-year), where actual momentum is strong, and too much at long horizons (five-year), where home price growth tends to mean revert. Case et al. (2012) and De Stefani (2021) document similar patterns in the data. One interpretation of this evidence is as a manifestation of extrapolation bias in the housing market context. The behavioral model of Glaeser and Nathanson (2017) generates this under-extrapolation in the short run and over-extrapolation in the long-run when agents neglect that past prices contain information about the beliefs of other agents and assume that other agents made the same mistake in the past.

This paper studies the question of expectation formation in the housing market with new data and new

[1] Bailey et al. (2018) show that house price expectations affect home purchase decisions, while Bailey et al. (2019) show that they affect mortgage leverage choice and mortgage rate. Favilukis, Ludvigson and Van Nieuwerburgh (2017); Kaplan, Mitman and Violante (2020); Greenwald and Guren (2021) debate the role of expectations and credit constraints in accounting for the boom and bust in U.S. house prices.

[2] Evidence that house price growth expectations are positively correlated with past realized house price growth goes back to Case and Shiller (1988). Recent contributions include Case et al. (2012); Fuster, Laibson and Mendel (2010); Greenwood and Shleifer (2014); Barberis et al. (2015); Liu and Palmer (2021); Giglio et al. (2021*a*,*b*).

tools. We use a unique data set of professional forecasters collected by the forecasting firm Pulsenomics. About 100 professionals from industry, policy, and academia predict house price growth in the current calendar year and in each of the next four (or five) calendar years.[3] The survey is quarterly and spans the period from 2010.Q1 until 2022.Q4 (52 quarters). What is unique about this data is that (i) it is a panel data set, i.e., we can track the individual forecasters over time, and (ii) that it contains a term structure of expectations. Existing survey data, including outside of the housing domain, typically contain no panel dimension and no or limited forecast data across longer horizons. We will show that both the panel and the term structure dimensions are important for inference.

While professionals may exhibit different expectations from individuals, their belief formation is of independent interest given (i) their expertise, (ii) the prevalence of professional survey data in economics and finance, and (iii) the fact that many individuals rely on professionals. Indeed, Fuster et al. (2022) show that, when households are given the opportunity to pick among different sources of information to help predict future house price changes, about half chose the forecast of housing experts.

We begin by documenting several previously-known and new empirical facts about house price expectations in our data set. First, like in the household survey data, there is substantial time series variation in average expectations that is correlated with past house price growth, but less than in the realized time series.

Second, there is substantial cross-sectional dispersion in forecasts in each survey. Using hand collected data from LinkedIn, we obtain a large set of demographic characteristics for each forecaster, including age, location, and detailed education and professional experience history. These individual characteristics explain an additional 10% points of panel variation in forecasts relative to a model with time fixed effects. Years of experience is the most important covariate among the individual characteristics.

Third, a new fact to the housing expectation literature is that the dispersion in forecasts is higher for short-run than for long-run forecasts.

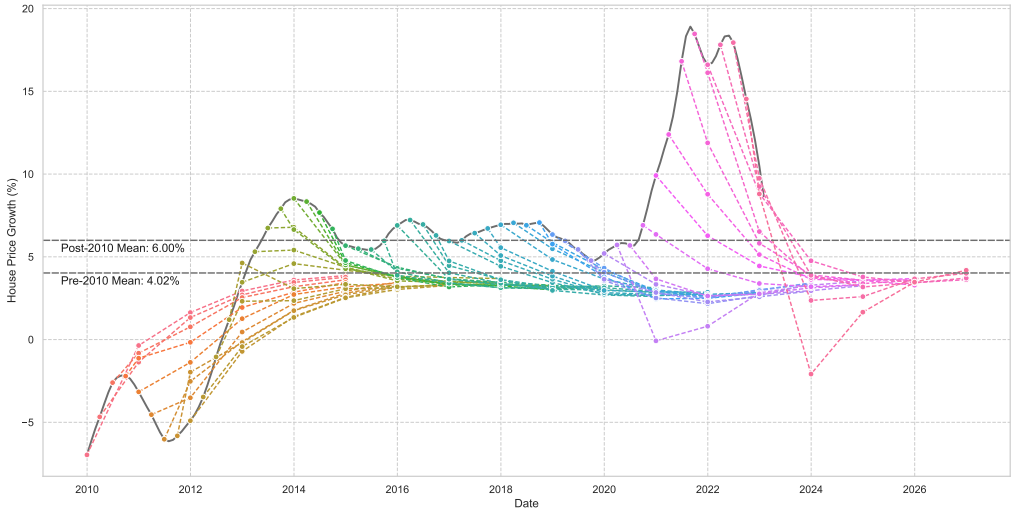
Fourth, we focus on the slope of a regression of forecast errors on the forecast revision. Coibion and Gorodnichenko (2012) interprets the negative sign and magnitude of the CG slope coefficient as evidence for limited- and sticky-information models, and specifically for over-reaction of the forecast to news. For house price growth, we also find a negative slope using individual forecaster data (panel regression). Using consensus data, in contrast, the slope turns positive. A similar sign reversal for individual versus consensus forecast data has been observed in other macro-economic and financial series by Bordalo et al. (2020), who propose diagnostic expectations to account for the pattern. Ours is the first paper to study individual and consensus forecasts jointly in the housing context. We add to the literature new facts about how the CG

[3]The Wall Street Journal runs a similar survey, but we are not aware of any analysis like ours with that data set.

slope coefficients change with the horizon of the forecast, and show that the horizon dependence is quite revealing for models of expectation formation.

Fifth, we find strong evidence that long-horizon forecasts are fully mean-reverting. Figure 1 shows the consensus forecast (i.e., averaged across forecasters) for each horizon in each quarter in the colored dashed lines. The solid black line shows the realized house price growth. Forecasters' prediction for annual house price growth four years from now is close to the unconditional mean of house price growth, indicating that the typical forecaster expects full mean-reversion by year four. This suggests that professional forecasters over-extrapolate over long-horizons, in contrast with households who the literature argued under-extrapolate over long horizons.[4] However, the literature hitherto does not have much data (and no panel data) on long-term forecasts of house price growth. And the statement of over- or under-extrapolation requires a model. As we will argue below, this evidence is consistent with a rational learning model.

Figure 1: Zillow/Pulsenomics Survey Forecasts of House Price Growth



Notes: The black solid line shows the realized annual house price growth. Each short colored dashed line with six small dots represents the actual annual house price growth at that quarter (the first circle), the average survey forecasts of house price growth of the current year (the second circle), and the average survey forecasts of house price growth in the following 4 years (subsequent four circles).

Sixth, we study the slope of a panel regression of the forecast on the realized value of house price growth. This tests the Mincer and Zarnowitz (1969) prediction that the rational forecast and the truth should on average move one-for-one with one another. In the data, we find very weak correlations between forecasts and realizations.

[4] In other words, households underestimate the long-term mean reversion, while professionals overestimate the long-term mean reversion.

Seventh, we look at the size of the forecast error, finding that the predicted house price growth in our sample is significantly below realized house price growth.

All these features of the data are usually interpreted as *prima facie* evidence against the rational expectation model.

The main finding of this paper is that a state-of-the-art Bayesian learning model with learning about the hidden state of the housing market and about model parameters goes a long way towards accounting for these features.

Similar to Bordalo et al. (2020); Farmer, Nakamura and Steinsson (2023), we model house price growth as containing a long-run mean, a latent mean-reverting component, and an i.i.d. noise term. The key deviation from rational expectations is that agents do not know the parameters of this data generating process, unlike in Bordalo et al. (2020), which assume agents know the value of the long-run mean and the persistence parameter governing the mean-reverting component. Learning about the long-run mean turns out to be particularly important for the results.

Using Bayes Law, our forecasters combine their priors with data on house price realizations and private signals to learn about the parameters and the state of the housing market. Unlike Farmer, Nakamura and Steinsson (2023), we allow for heterogeneity across forecasters. We allow for heterogeneity in the priors about the long-run mean of house price growth. We also allow for heterogeneity in signals. Each forecaster's private signal is a noisy version of house price growth that is i.i.d. across forecasters. The signal reflects the external data that each forecaster brings to bear on the prediction problem. The dispersion in the signal captures the heterogeneity in that information set. The signal dispersion can also capture differential limits on attention or information processing capacity across forecasters (e.g., Sims, 2003; Woodford, 2009; Van Nieuwerburgh and Veldkamp, 2009).

Even though the true parameters of the house price process are constant, the posterior (subjective) beliefs about these parameters are time-varying and heterogeneous because forecasters see different signals (Patton and Timmermann, 2010; Collin-Dufresne, Johannes and Lochstoer, 2016) and start from different priors about long-run mean house price growth. Cross-sectional dispersion in forecasts arises from prior and signal heterogeneity.

We estimate the parameters of the model using simulated method of moments with the help of cloud computing.

This model matches most, but not all aforementioned features of the data. It produces large in-sample forecast biases, negative CG regression slopes at the individual level and especially at longer horizons, and positive slopes at the aggregate level, realistic dependencies of house price growth forecasts on both contemporaneous and lagged realized house price growth, and substantial cross-sectional dispersion in

forecasts especially at longer horizons.

A key mechanism behind the model's success lies in the fact that the long-run mean house price growth rate is difficult to learn about. This is reflected by substantial uncertainty about the long-run mean. When forecasters have more diffuse priors, they assign more weight to their signals when making forecasts. The resulting forecasts appear excessively volatile or "over-reacting" to news about house price growth. Confronted with output from this model on a panel of forecasters, an econometrician would observe negative individual CG coefficients and a positive response of the forecast to past house price growth.

Generating substantial forecast dispersion requires both heterogeneity in priors and in signals. Signal heterogeneity helps but ultimately has limited bite in a world where forecasters are aware that their signals contain substantial noise. Heterogeneity in forecasters' prior beliefs about the long-run mean has a first-order impact on forecast dispersion, as long as uncertainty about the long-run mean is not too high. If uncertainty is too high, forecasters assign little weight to their priors when forming posteriors and heterogeneity in the prior fades away.

Taken together, the comprehensive set of moments we consider identify a set of model parameters that point to significant uncertainty and heterogeneity in the prior long-run mean house price growth rate. They also point to the importance of signal heterogeneity.

Where the learning model falls short is in generating a negative enough CG coefficient at the 1-year forecast horizon and enough forecast dispersion at the 1-year horizon.

To remedy this shortcoming, we embed two seminal concepts from behavioral economics—over-confidence and diagnostic expectations—in our learning model. An overconfident forecaster has a subjective signal precision that exceeds the objective one. Following Bordalo et al. (2020), a forecaster with diagnostic expectations forms distorted estimates for the mean-reverting state of the housing market. We show that either overconfidence or diagnostic expectations are tremendously helpful for matching the one-year individual GC slopes and one-year forecast dispersion. In contrast, the effect on longer-horizon forecasts from these behavioral frictions is small. Learning about the long-run mean is what helps the model match the longer-horizon CG slope and forecast dispersion. Combined, we are able to account for the full set of expectations moments, and arrive at a compelling synthesis of how to understand the extent of rationality and disagreement in house price forecasts.

To provide additional intuition for these results and to highlight the uniqueness of our data, we consider two simpler models: one where there is no heterogeneity across forecasters and one where there is no information on long-horizon forecasts.

The former is the model that a researcher who only has access to consensus (i.e., average) forecast data would be forced to estimate. When trying to fit the "no dispersion" model to the same set of moments

(minus the forecast dispersion moments of course), we obtain a significantly worse fit. We also obtain substantially different prior parameter estimates, especially for the crucial parameter that governs the prior uncertainty about long-run mean growth. This suggests that the individual CG regression slopes, which require a panel dataset to estimate, are important to identify a plausible learning model.

The latter model estimation, which only uses one-year ahead forecast moments, illustrates the importance of having access to a term structure of forecasts. We find that the fit deteriorates enormously, and the parameter estimates change substantially. Specifically, we estimate a lower prior mean for the persistence parameter with substantially more uncertainty around the estimate than in our benchmark model.

The rest of this paper is organized as follows. Section 2 introduces the Pulsenomics house price forecast data and documents key empirical facts on professional house price growth forecasts across time and forecasters. Section 3 sets up the model and the estimation algorithm. The results for the benchmark learning model are in Section 4. Section 5 adds overconfidence and diagnostic expectations to the benchmark model. Section 6 explains the importance of having a panel dimension and a term structure of forecasts in the data in order to identify the key parameters and account for the key moments of the survey data. Section 7 concludes. The Appendix contains additional empirical results, model derivations, and auxiliary model results.

2 Facts About Experts' House Price Growth Forecasts

Using a unique data set of professional forecasters collected by the forecasting firm Pulsenomics, we document several important empirical facts about house price growth forecasts, including time series variation, cross-sectional dispersion and heterogeneity, relation to past house price growth, and other forecast anomalies.

2.1 Data

The forecast data we use come from Zillow Home Price Expectations Survey. The survey is conducted by Pulsenomics, an independent research and index product development firm. The sample period of our survey data is 2010Q1 to 2022Q4 (52 quarters). Every quarter, more than 100 professionals are surveyed to predict the year-by-year home price growth over a five-year horizon. This diverse expert panel consists of professionals and economists from industry, policy, and academia. The survey for each quarter is conducted in the middle of the second month of that quarter, e.g., the 2010Q1 survey is conducted in the middle of Feb 2010.[5] The survey data is unique in the following ways. First, it is a panel data set. Each individual

[5]More details about the timing of the survey and the data availability can be found in Appendix A.

forecaster can be tracked over time. This allows us to study belief formation at both the individual and the consensus levels. Second, it contains a term structure of forecasts with relatively long horizons (from current year to 4 years ahead). Information from the term structure of expectations turns out to be very informative for estimating the belief formation process.

We augment this data set by manually collecting a large set of demographic characteristics for each forecaster from LinkedIn. It contains characteristics in the following categories:

- Personal: (i) Age, (ii) Region metropolitan area, non-metropolitan area, and outside the U.S. This variable is categorized based on the location of the forecaster.
- Educational background: (i) Highest degree (PhD, Master's, or Bachelor's), (ii) Major (classified as economics, STEM, or both).
- Professional experience: (i) Length of work experience (calculated as the number of years since the forecaster started her first job), (ii) Industry (classified based on the firm that currently employs the forecaster), (iii) Number of skills (calculated based on the skills that a forecaster reports on her LinkedIn page. Each reported skill can be endorsed by other people.), (iv) Skills related to real estate (based on self-reported skills).
- Number of followers: defined as the total number of followers on her LinkedIn page. This variable is a proxy for the forecaster's network of professional connections and sphere of influence.

The data we use for the house price index is the Zillow Home Value Index (ZHVI). This index is calculated by Zillow based on more than 100 million U.S. homes. Since forecasters in the Zillow Home Price Expectations Survey are asked to predict the changes in the ZHVI, we naturally use this index as our data source for the house price index. The index is available monthly and data for a given month is published on the third Thursday of the following month. This data series starts in 1996.

2.2 Time-series and Cross-section of Forecasts

Expectations of house price growth vary over time and have large dispersion across individuals. Figure 2 shows the time series of short-term and long-term expected house price growth, alongside past realized house price growth. Consistent with the findings from household surveys, there is significant time-series variation in professional house price growth forecasts. However, the magnitude of this time-series variation (right axis, solid line) is smaller than the time-series variation in the realized house price growth (left axis, dashed line). Another important pattern is that the variation in long-term expectations (plotted in panel b)

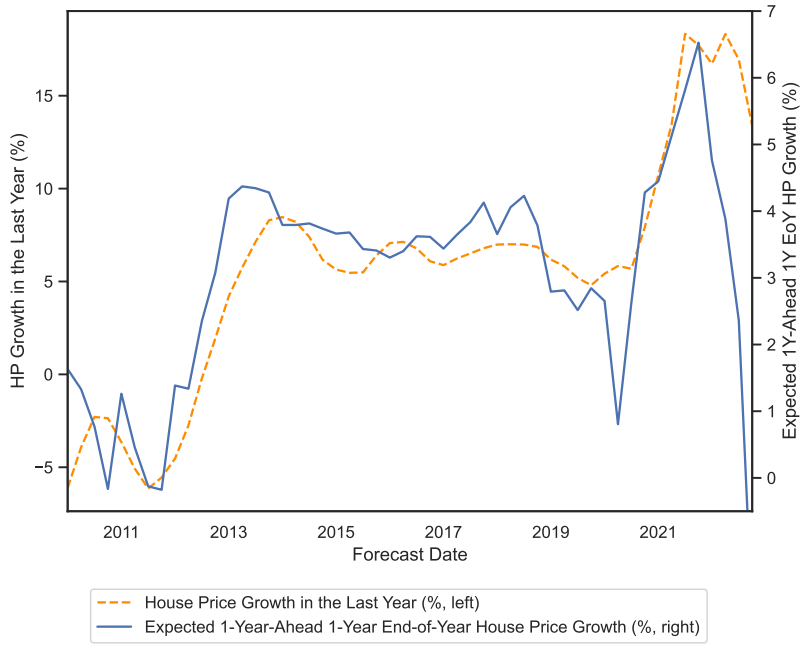
is smaller than the variation in short-term expectations (panel a). This finding is consistent with Figure 1, which showed that long-term forecasts fluctuate only modestly around the long-run mean.

Figure 3 shows the cross-sectional dispersion of house price forecasts in the 2022.Q4 survey. There is substantial dispersion across individuals. The cross-sectional variation in forecasts exceeds the time-series variation in average forecasts of house price growth. For example, some professionals forecast a house price growth of 5% in 2023 while others expect a decrease of 20% (panel a). The forecast for 2026 still shows wide dispersion (panel b), but is smaller than the dispersion about 2023. This large cross-sectional dispersion, ranging from 0% to 10% in 2022.Q4, stands in contrast with the modest fluctuations in the time series of average four-year ahead house price growth forecasts, which move around in a much smaller range of around 2.5% to 4.5%.

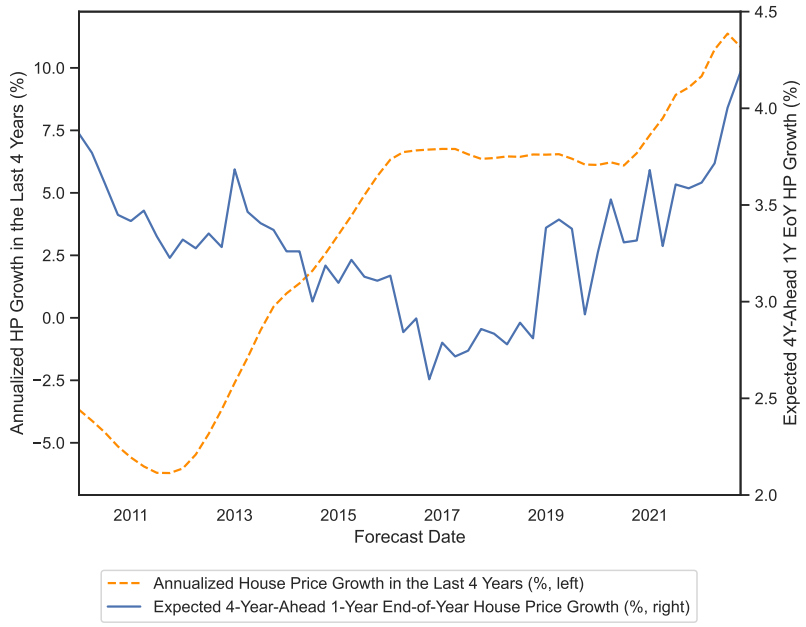
Table 1 explores the role of individual characteristics in accounting for the cross-sectional variation in house price forecasts. It estimates a panel regression of individual house price growth forecasts on forecaster characteristics. To provide a baseline, Column (1) regresses the forecasts on a time fixed effect. This picks up all variation in past house price growth as well as in all other macro variables. Column (2) adds the forecaster characteristics, introduced above. The R^2 increases from 29.5% to 39.7% for the one-year ahead forecast and from 3.7% to 11.1% for the 4-year ahead forecast. The most significant variable is length of experience. Forecasters with 15-25 years of professional experience predict quarterly house price growth over the next year (four-years ahead) that is nearly 1% (0.6%) point lower than forecasters with more than 35 years of experience.

Despite the rich set of covariates, the table makes clear that there is a lot of unexplained heterogeneity in these forecasts. This leaves open the door for unobservable heterogeneous information across forecasters.

Figure 2: Time-Series of House Price Growth Forecasts



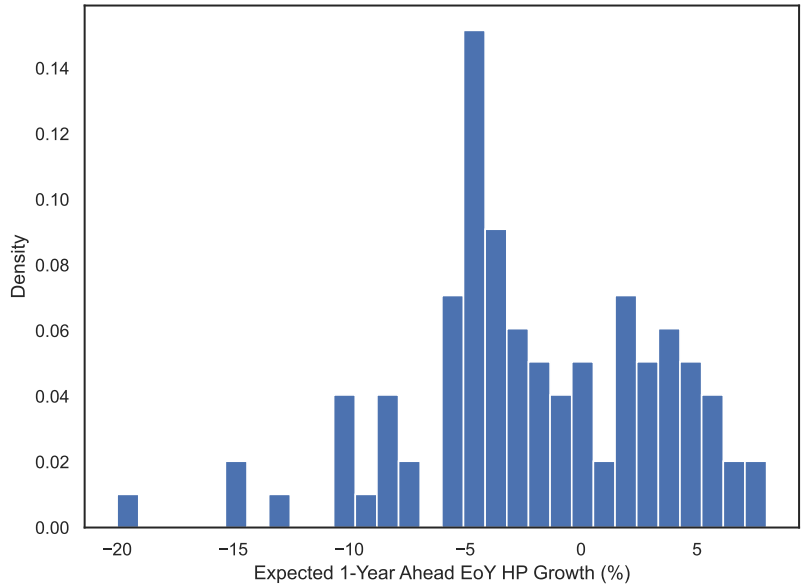
(a) 1-Year Ahead



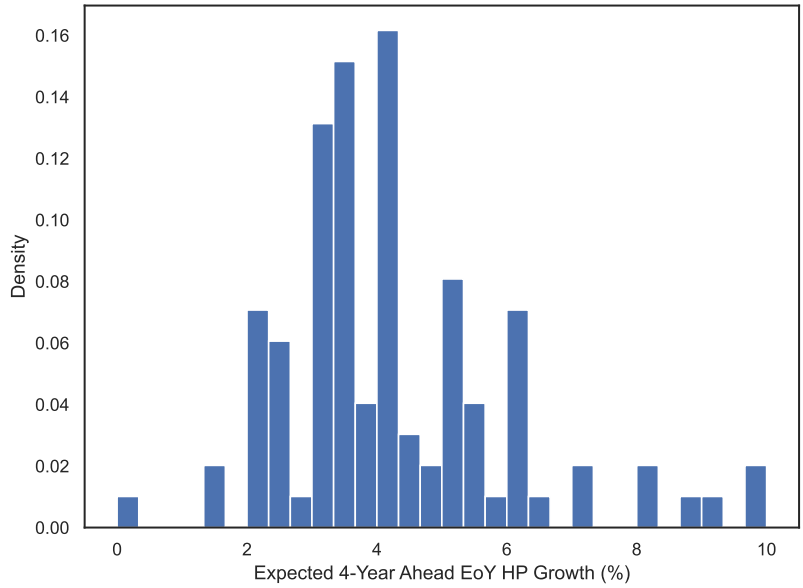
(b) 4-Year Ahead

Notes: This figure presents the time series of average professional forecasts of U.S. house price growth and past house price growth. Panel A shows the expected 1-year-ahead 1-year end-of-year house price growth (e.g., from Dec/31/2010 to Dec/31/2011 for 2010Q1 survey, solid line, right y-axis) and the actual house price growth in the last year (e.g., from Dec/31/2008 to Dec/31/2009 for 2010Q1 survey, dashed, left y-axis). Panel B shows the expected 4-year-ahead 1-year end-of-year house price growth (e.g., from Dec/31/2013 to Dec/31/2014 for 2010Q1 survey, solid line, right y-axis) and the actual house price growth in the last year (e.g., from Dec/31/2008 to Dec/31/2009 for 2010Q1 survey, dashed, left y-axis).

Figure 3: Cross-Sectional Dispersion of House Price Growth Forecasts



(a) 1-Year Ahead



(b) 4-Year Ahead

Notes: This figure shows the cross-sectional dispersion of house price growth forecasts in 2022Q4 survey. Panel A shows the histogram of expected 1-year-ahead 1-year end-of-year house price growth (from Dec/31/2022 to Dec/31/2023). Panel B shows the histogram of expected 4-year-ahead 1-year end-of-year house price growth (from Dec/31/2025 to Dec/31/2026).

Table 1: Demographic Characteristics and House Price Growth Forecasts

	1Y EoY HP Growth (In 1 Year)		1Y EoY HP Growth (In 4 Years)	
	(1)	(2)	(3)	(4)
Length of Experience (yrs.)				
>35		(omitted)		(omitted)
25-35		-1.004***		-0.427*
		(0.316)		(0.252)
15-25		-0.919***		-0.569**
		(0.290)		(0.237)
<15		-0.533*		-0.598***
		(0.272)		(0.219)
Highest Degree				
Bachelor		(omitted)		(omitted)
Master		-0.101		0.255
		(0.374)		(0.398)
Doctor		0.297		0.646
		(0.354)		(0.424)
Region				
Metro Area		(omitted)		(omitted)
Non-Metro Area		-0.098		-0.077
		(0.258)		(0.203)
Foreign		0.190		0.352
		(0.345)		(0.266)
Major				
Econ		0.346		-0.152
		(0.233)		(0.223)
STEM		-0.688		-0.332
		(0.485)		(0.253)
Econ * STEM		0.565		0.368
		(0.564)		(0.419)
Skill Number				
<25% Percentile		(omitted)		(omitted)
25% - 50% Percentile		-0.108		0.233
		(0.320)		(0.276)
50 - 75% Percentile		-0.199		0.040
		(0.269)		(0.183)
>75% Percentile		-0.571**		0.096
		(0.283)		(0.268)
Has Skills Related to Real Estate		0.016		0.010
		(0.235)		(0.225)
Number of Followers				
<25% Percentile		(omitted)		(omitted)
25% - 50% Percentile		0.565		0.310
		(0.374)		(0.272)
50 - 75% Percentile		0.507		0.098
		(0.360)		(0.315)
>75% Percentile		0.323		0.125
		(0.372)		(0.258)
Survey Quarter FE	Y	Y	Y	Y
R-Squared	0.295	0.397	0.037	0.111
Number of Observations	5,503	3,052	5,503	3,052

2.3 Forecast Anomalies

As a motivation for the model, we document several key empirical results. These empirical patterns are typically interpreted as evidence that expectations depart from rationality.

First, we show that house price growth expectations exhibit correlations with past returns. We estimate the sensitivity coefficients from:

$$E_{it}[y_{t+h}] = a + by_{t-j} + \varepsilon_{t+h}, \tag{1}$$

where y_{t+h} denotes one-year house price growth h years from now, y_{t-j} is lagged annualized house price growth over the past j years, and $E_{it}[y_{t+h}]$ the individual forecast of that future house price growth. Panel A of Table 2 shows the results. Specifically, short-term forecasts are strongly positively correlated with past realized growth, while long-term forecasts are insensitive to past realized growth.

Panel B contrasts these correlations of expected house price growth with the same correlations for realized house price growth, estimated for the same period 2010–2022. Realized house price growth shows strong serial correlation that diminishes with the horizon. Professional forecasters "under-extrapolate" at all future horizons $h = 1, 2, 3, 4$. Put differently, they over-estimate the degree of mean reversion in house prices growth.

Second, individual house price forecasts have persistent biases, $e_{i,t+h|t} \equiv y_{t+h} - E_{it}[y_{t+h}]$. The average bias in the panel data set of forecasters can be estimated from the regression:

$$e_{i,t+h|t} = a^{fb} + \varepsilon_{i,t+h}, \tag{2}$$

Panel A of Table 3 shows that forecasts are systematically lower than actual house price growth during our 2010–2022 sample, so that the bias a is large and positive.

Third, we estimate (inverse) Mincer-Zarnowitz regressions of forecasts on contemporaneous realizations:[6]

$$E_{it}[y_{t+h}] = a^{mz} + b^{mz}y_{t+h} + \varepsilon_{i,t+h}, \tag{3}$$

Full information rational expectations implies that the truth and the forecast move one-for-one, i.e., $b^{mz} = 1$ in the above regression. We find that forecasts move less than one-for-one with the truth. Panel B of Table 3 shows slope estimates close to zero for $h = 1$ and even negative for $h = 2, 3, 4$.

[6]The standard Mincer-Zarnowitz regression is defined as: $y_{t+h} = \widetilde{a}^{mz} + \widetilde{b}^{mz}E_{it}[y_{t+h}] + \widetilde{\varepsilon}_{i,t+h}$. However, in a pooled individual-level panel regression, the dependent variable y_{t+h} is the same for every individual in a single cross-section, which can cause the regressors to correlate with error terms. To avoid a biased estimation of the coefficient, we use an inverse version of the Mincer-Zarnowitz regressions. In these inverse MZ regressions, full information rational expectation still implies that $b^{mz} = 1$, meaning that realized values move one-for-one with forecasts. If $b^{mz} \neq 1$, the interpretation of the coefficient can be adjusted accordingly.

Table 2: Effect of House Price Growth History on House Price Growth Forecasts

	Panel A: Expected 1Y HP Growth at Different Horizon:									
	Current Year		**1-Year Ahead**		**2-Years Ahead**		**3-Years Ahead**		**4-Years Ahead**	
Past 1Y HPG	0.549***		0.165***		0.044**		0.015		0.003	
	(0.040)		(0.042)		(0.017)		(0.011)		(0.011)	
Past 4Y HPG		0.480***		0.131**		0.019		-0.003		-0.009
		(0.061)		(0.050)		(0.022)		(0.014)		(0.013)
Num. of Obs.	5,503	5,503	5,503	5,503	5,503	5,503	5,503	5,503	5,503	5,503
R-Squared	0.628	0.410	0.116	0.063	0.014	0.002	0.002	0.000	0.000	0.001
	Panel B: Realized 1Y HP Growth (2010-2022) at Different Horizon:									
	Current Year		**1-Year Ahead**		**2-Years Ahead**		**3-Years Ahead**		**4-Years Ahead**	
Past 1Y HPG	0.735***		0.464**		0.194		0.169		0.269**	
	(0.115)		(0.187)		(0.132)		(0.141)		(0.127)	
Past 4Y HPG		0.705***		0.498**		0.291*		0.291*		0.437***
		(0.167)		(0.215)		(0.157)		(0.152)		(0.142)
Num. of Obs.	52	52	48	48	44	44	40	40	36	36
R-Squared	0.739	0.560	0.283	0.297	0.0722	0.191	0.0642	0.211	0.158	0.425

Notes: Panel A shows the relationship between house price forecasts and past house price growth. The coefficients are estimated from individual-level panel regressions. Panel B shows regression estimates of actual house price growth dependence on past growth. All past house price growth variables are annualized. Standard errors are in parentheses. For time-series regressions, standard errors are Newey-West with lag length of $L = \left\lceil 0.75 \times T^{1/3} \right\rceil$. For individual-level panel regressions, standard errors are clustered by both forecaster and time. ***: p<0.001; **: p < 0.01; *: p<0.05.

Fourth, forecast revisions negatively predict future forecast errors. We estimate the Coibion-Gorodnichenko regressions:

$$e_{t+h|t} = a^{cg} + b^{cg} \left(E_{it}[y_{t+h}] - E_{it-1}[y_{t+h}] \right) + \varepsilon_{t+h}, \tag{4}$$

These regressions test the null of full-information rational expectations ($b^{cg} = 0$) against the alternative of limited- or sticky-information models. If $b^{cg} > 0$, the forecasts suffer from "under-reaction," in the sense that an increase in the forecast leaves the new forecast still too low on average. If $b^{cg} < 0$, the forecasts suffer from "over-reaction." Panel C of Table 3 finds that professional forecast errors are negatively correlated with forecast revisions, and more so at longer horizons. Professional house price forecasts suffer from over-reaction.

Table 3: Other Forecast Anomalies

	Forecast Horizon				
	Current Year	**1-Year Ahead**	**2-Years Ahead**	**3-Years Ahead**	**4-Years Ahead**
Panel A: Bias					
a	1.724***	3.593***	4.935***	5.240***	5.072***
	(0.387)	(0.630)	(0.565)	(0.574)	(0.628)
Num. of Obs.	5,503	5,117	4,737	4,352	3,911
Panel B: Inverse MZ Regressions					
b	0.596***	0.124**	-0.043**	-0.056***	-0.064***
	(0.058)	(0.049)	(0.021)	(0.017)	(0.016)
Num. of Obs.	5,503	5,117	4,737	4,352	3,911
R-squared	0.548	0.048	0.005	0.010	0.016
Panel C: CG Regressions					
b	-0.201	-0.219	-0.514***	-0.574***	-0.418***
	(0.144)	(0.258)	(0.070)	(0.007)	(0.062)
Num. of Obs.	4,728	4,407	4,065	3,732	2,589
R-squared	0.019	0.006	0.029	0.032	0.017

Notes: This table shows the regression estimates of other forecast anomalies at different horizons. Panel A shows the estimates of bias in equation 2. Panel B shows the estimates of inverse MZ regression specified in equation 3. Panel C shows the estimates of CG regression in equation 4. All coefficients are estimated from individual-level panel regressions. Standard errors are in parentheses. Standard errors are clustered by both forecaster and time. ***: p<0.001; **: p < 0.01; *: p<0.05.

3 Model

We develop a model to understand these key empirical results. The model aims to simultaneously match the panel of individual house price forecasts and the heterogeneity across forecasters. It provides a laboratory to quantitatively assess the importance of imperfect information, learning, and behavioral biases in shaping expectations.

3.1 Setup

The key assumption of the model is that forecasters do not know and must learn the underlying state of the housing market and the parameters that govern the data-generating-process (D.G.P.) of the house price growth process. Our setup is similar to Bordalo et al. (2020) and Farmer, Nakamura and Steinsson (2023) who study macroeconomic and financial time series. Unlike Bordalo et al. (2020), who also assume the presence of a latent state, we assume that forecasters additionally do not know the underlying parameters of the DGP. They need to estimate the parameters in order to forecast house price growth. Unlike Farmer, Nakamura and Steinsson (2023), we additionally account for forecasts at multiple horizons and for disagreement among forecasters by allowing different forecasters to receive heterogeneous noisy signals. Farmer, Nakamura and Steinsson (2023) work with consensus forecasts and a single-period forecast horizon.

House Price Growth Process We assume that the house price growth y_t, measured by log changes in the ZHVI, contains a long-term mean, $\bar{\mu}$, a mean-reverting component x_t, and a noise term e_t:

$$y_t = \overline{\mu} + x_t + \sqrt{\overline{\gamma}\overline{\sigma}^2} e_t \quad e_t \sim N(0,1) \tag{5}$$

$$x_t = \overline{\rho} x_{t-1} + \sqrt{(1-\overline{\gamma})\,\overline{\sigma}^2}\,\eta_t \quad \eta_t \sim N(0,1) \tag{6}$$

where $\overline{\gamma}$ governs the proportion of variance in total house price growth, $\overline{\sigma}^2$, that comes from the noise component e_t. The noise e_t term captures the fact that aggregate house price growth data from Zillow is a noisy measure of the true house price growth. Indeed, Zillow revises its own historical past realized price growth as new data becomes available. Through the persistent state variable x, the model captures the fact that house price growth is serially correlated, unlike price growth in many other financial time series, and may exhibit mean-reversion. The model is similar to "the kernel of truth" laid out in Bordalo et al. (2020), except for the presence of a long-term mean $\bar{\mu}$. The lines over the parameters indicate that these are the true parameters that govern the objective D.G.P., unknown to the forecasters.

Forecasters' Heterogeneous Signal Extraction Process Each quarter, the Pulsenomics Survey asks its 100 participants to make forecasts for house price growth over the current calendar year and each of the next three calendar years. Let T denote the current quarter during which forecasters are asked to fill out the survey. Each forecaster observes (is shown) the realized price growth up to previous quarter $(T-1)$ when making forecasts for future house price growth, but does not know the house price in the current quarter T. We formalize the timing of the survey in Appendix A.

Each forecaster i at time T has information set $s_{i,T}$, defined as

$$S_{i,T} := (s_{1,t}, \ldots, s_{1,T-1}, s_{i,T})' \tag{7}$$

The information set contains all historical realized house price growth and a private signal about current and future house price growth, or

$$s_{i,t} = \begin{cases} y_t & t \leqslant T-1 \\ y_T + \alpha_{i,T}, \quad \alpha_{i,T} \sim N(0, \overline{\Gamma}^2 \overline{\sigma}^2) & t = T \end{cases} \tag{8}$$

Notice that $\overline{\Gamma}$ denotes *objectively* how much noise (variance of $\alpha_{i,T}$) is in forecasters' information signal. The noise $\alpha_{i,T}$ are assumed to be i.i.d. across investors and time. $\overline{\Gamma}$ differs from $\overline{\gamma}$ in that the latter captures the proportion of the noise in the realized house price growth, the part that attributes to the one-period shock not going into neither x_t nor μ.

The parameter Γ (without the $\bar{\cdot}$) captures forecasters' subjective beliefs about the amount of noise in their signals. In the case of $\Gamma = \overline{\Gamma}$, which is our benchmark case, forecasters correctly understand the true amount of noise that their private signals contain.

Ultimately, their different information sets $s_{i,T}$ lead them disagree about the current and future house price growth, which should be reflected in their house price forecasts. This heterogeneity in forecasters' signals is rooted in a literature on investors' noisy signals and rational attention (e.g., Sims, 2003; Woodford, 2009; Van Nieuwerburgh and Veldkamp, 2009). One can interpret the differences in signals as different forecasters using different predictors to estimate house price growth or paying different amount of attention to the latest developments in housing markets.

Forecasters' Subjective Expectation Formation Process Based on past realized house price growth and their own individual signals, the forecasters make inference about the values of the parameters and the latent state of the housing markets. They correctly understand the underlying structure of the D.G.P. in Equation (5) and (6), and the presence of noise in their signals as in Equation (8). However, they do not

know the true values of the underlying parameters and states. They need to make inference about the parameters ($\overline{\mu}$, $\overline{\rho}$, $\overline{\sigma}$, $\overline{\gamma}$) and the latent state $\{x_t\}_{t=0}^{T}$ in order to form expectations about future house prices based on their information set $S_{i,T}$. They are Bayesian learners. We assume in our benchmark case that the forecasters know the value of $\overline{\Gamma}$, the objective signal volatility.

From the perspective of the forecasters, their learning process takes the following form:

$$s_{i,t} = \mu_{i,t} + \widetilde{x}_{i,t} + \widetilde{e}_{i,t}, \qquad \widetilde{e}_{i,t} \sim N(0, H_{i,t}) \tag{9}$$

$$\widetilde{x}_{i,t} = \rho_{i,t}\widetilde{x}_{i,t-1} + \sqrt{1-\gamma_{i,t}}\sigma_{i,t}\widetilde{\eta}_{i,t}, \qquad \widetilde{\eta}_{i,t} \sim N(0,1) \tag{10}$$

where

$$H_{i,t} = \begin{cases} \gamma_{i,t}\sigma_{i,t}^2, & t \leqslant T-1 \\ (\gamma_{i,t} + \Gamma^2)\sigma_{i,t}^2, & t = T \end{cases} \tag{11}$$

Because of the heterogeneity in forecaster information, the parameter estimates ($\mu_{i,t}$, $\rho_{i,t}$, $\sigma_{i,t}$, $\gamma_{i,t}$) are heterogeneous across forecasters, as indicated by the subscript "i" (Patton and Timmermann, 2010). The state estimates, $\widetilde{x}_{i,t}$, are also heterogeneous across forecasts and subjective, as indicated by " $\widetilde{\ }$ ". The disagreement across forecasters is bounded as all of the forecasters observe the same history of realized house price growth up until the previous quarter, or $s_{i,t} = y_t, \forall t = 1, ..., T-1$.

Due to parameter uncertainty and learning (Collin-Dufresne, Johannes and Lochstoer, 2016), the parameter estimates ($\mu_{i,t}$, $\rho_{i,t}$, $\sigma_{i,t}$, $\gamma_{i,t}$) are time-varying, because they are the posterior mean estimates given the realized data up until that point, signals, and prior distributions.

Forecasters' Prior Beliefs We assume the following prior distributions for the parameters:

$$\mu_0 \sim N(\mu_\mu, \sigma_\mu^2),$$
$$\rho_0 \sim N(\mu_\rho, \sigma_\rho^2),$$
$$\sigma_0^2 \sim IG(A, B),$$
$$\gamma_0 \sim Beta(C, D)$$

These priors are common across forecasters, except that we introduce an additional parameter λ which captures heterogeneity across forecasters in the prior belief about the long-run mean HPA. Forecasters prior mean house price growth beliefs are uniformly distributed over the range $[\mu_\mu - \frac{1}{2}\lambda, \mu_\mu + \frac{1}{2}\lambda]$.

The agents in the model know the parameters that govern these prior distributions $(\mu_\mu, \sigma_\mu^2, \lambda, \mu_\rho, \sigma_\rho^2, A, B, C, D)$.

We start agents off in 1996 with these priors.

Upon receiving new signals, forecasters update their prior beliefs to arrive at posterior beliefs for the parameters ($\mu_{i,t}$, $\rho_{i,t}$, $\sigma_{i,t}$, $\gamma_{i,t}$) using Bayes Rule. Given these new estimates for parameters, they estimate the underlying mean-reverting state $\widetilde{x}_{i,t}$ using Kalman Filtering and Smoothing.

We derive analytical expressions for these posterior means based on marginal distributions of these parameters. We then perform Gibbs-Sampling to combine the marginal posteriors to arrive at the posterior mean of the forecasts for future house price growth. Appendix B contains the details.

Given their posterior parameter and state estimates, the forecasters form their forecasts for the m-quarter ahead house price growth:

$$E_{i,t}(y_{t+m,1}) = \mu_{i,t} + \rho_{i,t}^{m}\widetilde{x}_{i,t}, \tag{12}$$

where $\mu_{i,t}$ is the long-run mean forecast and $\widetilde{x}_{i,t}$ is the transitory component that affects the short- and medium-term forecast. Appendix A details how we compute the annual forecasts that mimic the exact structure and timing of the Pulsenomics survey.

Incorporating Behavioral Biases to the Expectation Formation Process Following a literature on over-confidence, we introduce an additional parameter $\phi := \overline{\Gamma}/\Gamma$. When $\phi > 1$, the objective signal noise is larger than the subjective signal noise. Forecasters' perceived signal precision is greater than the actual signal precision, i.e. they are overconfident about the precision of their signals. Conversely, when $\phi < 1$, forecasters are under-confident. The benchmark model sets $\phi = 1$. We incorporate ϕ into our learning model and examine its impact on the model's fit to the data in Section 5.

Following the literature on diagnostic expectations (Bordalo et al., 2020), we allow forecasters to have potentially distorted estimates of the state of the housing market, $\widehat{x}_{i,t}$, when applying the Kalman filter, given their posterior estimates for the parameters. The degree of distortion is governed by the parameter θ, with a positive (negative) θ leading to an overreaction (underreaction) to news. When $\theta = 0$, forecasters apply the optimal Kalman gain and we recover our benchmark rational learning model. We discuss this extensionin Section 5, and provide the technical details on how to incorporate diagnostic expectations in our learning model in Appendix C.

3.2 Benchmark Model Estimation

We, the econometricians, do not know the parameters of the prior distributions and must estimate them. We estimate the six parameters, $\delta = (\mu_\mu, \sigma^2_\mu, \mu_\rho, \sigma^2_\rho, \Gamma, \lambda)$ that govern our benchmark rational learning model using Simulated Method of Moments (SMM). We fix the parameters A, B, C, D that govern the prior distributions of σ^2_0 and γ. In extensions, we estimate additional parameters ϕ and θ to evaluate the impact of behavioral biases in forecasts. These parameters are set to one and zero, respectively, for our baseline analysis.

Our benchmark estimation aims to match salient features of the individual forecasters' panel data set as well as forecast dispersion, 36 moments in all. The technical details of the estimation are relegated to Appendix D. Given the computational complexities, the model is estimated on a supercomputer using Microsoft Azure.

Given a candidate parameter vector δ, we proceed as follows

1. Obtain N forecasters' heterogeneous signals by
   (a) simulating normally distributed private signals $\alpha_{i,t} \sim N(0, \bar{\Gamma}^2\bar{\sigma}^2), \quad i = 1, ..., N$
   (b) and adding them to the current quarter's realized price growth: $s_{i,t} := y_t + \alpha_{i,t}$.
2. Perform Bayesian learning for each forecaster i, using $s_{i,t}$ as inputs, and prior parameters $(\mu_\mu, \sigma^2_\mu, \mu_\rho, \sigma^2_\rho)$. This leads to individual forecasters' house price growth prediction, $E_{i,t}(y_{t+h})$, for different horizons h and periods $t = 1, ...T$.
3. Construct the 36 model-implied moments based on these forecasts $E_{i,t}(y_{t+h})$ as follows:
   (a) Estimate panel regressions based on individual forecasts $E_{i,t}(y_{t+h})$ and actual realized house price growth y_{t+h}. We use four forecasting horizons for each regression moment: $h = 1, 2, 3, 4$.
      i. Bias: Regress forecast error $y_{t+h} - E_{i,t}(y_{t+h})$ on a constant; collect mean estimates (1 moment per horizon).
      ii. Coibion-Gorodnichenko regression: Regress forecast error $y_{t+h} - E_{i,t}(y_{t+h})$ on a constant and the forecast revision $E_{i,t}(y_{t+h}) - E_{i,t-1}(y_{t+h})$; collect intercept and slope (2 moments per horizon).
      iii. Mincer-Zarnowitz regression: Regress forecast $E_{i,t}(y_{t+h})$ on a constant and the realization y_{t+h}; collect intercept and slope (2 moments per horizon).
      iv. Forecast on lagged realization: Regress forecast $E_{i,t}(y_{t+h})$ on constant and realized house price growth y_{t-j} for $j = 0.25, 2, 4$; collect slope (3 moments per horizon).

(b) Forecast dispersion for each $h = 1, 2, 3, 4$

i. We compute the cross-sectional variance of the individual forecasts in each quarter, and average over quarters:

$$Disp_h = \frac{1}{T}\sum_{t=1}^{T} Var_t(E_{i,t}(y_{t+m,h}))$$

4. Minimize the distance between the model-implied and the empirical moments through searching over the parameters δ:

$$\hat{\delta} = argmin_{\delta} \sum_{k=1}^{K} m_k(\delta)$$

where the k^{th} moment condition is the absolute difference between the empirical and the model-implied moment weighted by the standard error of that moment.The weighting scheme recognizes estimation uncertainty in the empirical moments, giving more weight to more precisely estimated moments.[7]

$$m_k(\delta) = \frac{|EmpiricalMoment_k - ModelImplied_k|}{SE(EmpiricalMoment_k)} \quad k = 1, ...K$$

We start our estimation process in 1996.Q1, when the ZHVI time series starts. This means that the prior distribution refers to 1996.Q1 and that the forecasters in our model start their learning process in 1996.Q1. The model produces quarterly model-implied house price growth forecasts from 1996.Q1 until 2022.Q4. Since Pulsenomics only started to collect experts' survey data from 2010.Q1 onwards, we use the model-implied expert forecasts from 2010.Q1 until 2022.Q4 to match the empirical moments.[8]

Given individual forecasts, we can compute the model-implied consensus forecast as the average forecast among the N forecasters in the model:

$$E_t^c(y_{t+h}) = \frac{1}{N}\sum_{i=1}^{N} E_{i,t}(y_{t+h})$$

We can also compute the moments in step 3(a) based on the consensus forecast rather than the individual forecast in both model and data. We return to this in Section 6.

[7] This is what we do for the 32 regression moments for which we can estimate the empirical standard error. We further increase the weights of 4 CG slope moments by 10 times due to their importance in understanding individual forecasters' expectations. Since we do not have a standard error for the forecast dispersion, for the 4 forecast dispersion moments, we use the absolute error scaled by the average of the weights of the 32 regression moments for which we have the empirical standard error.

[8] We also compare model-implied forecasts between 1996.Q1 and 2009.Q4 to realized house price growth; this is moment not targeted in the estimation.

4 Results Benchmark Learning Model

4.1 Parameter Estimates

Table 4 shows the estimated parameters in $\hat{\delta}$ that govern the prior distributions of μ and ρ, as well as the signal dispersion parameter Γ.

Table 4: Parameter Estimation for Different Estimation Procedures

	Estimation Procedure		
	Benchmark	**No LR.**	**No Disp.**
Prior About Long-run Mean			
μ_μ	0.762	0.621	0.810
σ_μ	0.218	0.188	0.040
Heterogeneity of Prior About Long-run Mean			
λ	2.928	2.335	
Prior About Persistent Parameter			
μ_ρ	0.659	0.538	0.685
σ_ρ	0.023	0.182	0.018
Information Heteorgeneity			
$\bar{\Gamma}$	1.274	0.977	

Notes: This table presents the estimates for the key parameters from our optimization process. The first column is for the benchmark model. The second column ("No LR.") excludes empirical moments based on the 2-year, 3-year and 4-year ahead forecasts in the optimization. The third column ("No Disp.") is for a simpler model that shuts down the heterogeneity in prior beliefs and signals and estimates the model by matching moments based on consensus forecasts. Appendix D provides more details on the optimization procedure.

Figure A1 plots the prior distributions of μ_0, ρ_0, and σ_0 implied by these parameter estimates. The prior mean long-run growth rate of μ_μ =0.762% per quarter is consistent with the historical sample average prior to 1996. Indeed, the Case-Shiller house price index data, which start in 1987, show a quarterly house price growth rate of 0.69% for the period 1987—1995, prior to the start of our estimation in 1996. The standard deviation of the prior is 0.218%, which reveals substantial prior uncertainty about the long-run mean house price growth rate.

The parameter $\lambda = 2.928\%$ captures heterogeneity in the prior mean μ_μ across forecasters. The data insist on a large amount of cross-sectional heterogeneity in beliefs about long-run HPA. As we shall see, learning about the long-run mean is very difficult (slow), so that much of the prior uncertainty persists over time and in the cross-section. Uncertainty about long-run mean house price growth will be a key feature that helps the learning model account for various anomalies.

The parameter ρ governs the persistence of the latent state of the housing market, x. We estimate a

tight prior distribution centered around a mean of 0.659% (quarterly value) with a standard deviation of 0.023%. The long-run moments we consider in our estimation help us to identify the ρ estimates. This is evidenced by the large difference between the ρ estimates of "Benchmark" and those of"No-LR", which does not include long-run moments.

Finally, we estimate a substantial amount of signal dispersion across forecasters. Recall that forecasters receive an i.i.d. draw of a signal about current quarterly house price growth, $\alpha_{i,T}$, which is normally distributed with mean zero and signal precision $(\overline{\sigma}\overline{\Gamma})^{-2}$. Our estimate for $\overline{\Gamma}$ of 1.274% indicates that a forecaster with one-standard deviation higher signal predicts a 1.274% point higher quarterly house price growth rate.

For simplicity, we fix the prior distribution of $\sigma_0^2 \sim IG(2.82, 3.27)$ rather than estimate it. The mean and volatility of the prior for house price growth volatility in 1996.Q1, the start date of the Zillow data, are set equal to the observed mean and volatility of annual house price growth from 1975 to 1995 from the Freddie Mac house price index data. We also fix γ at 0.002, which implies that 4% of total price growth volatility, on average, is noise.[9]

4.2 Model Fit

Table 5 column (1) shows the 36 empirically-observed moments that the model aims to fit. Column (2) reports the moments implied by the Benchmark model. Note that we only have 6 parameters to match these 36 moments, so that a close fit is far from guaranteed.

[9]This can be interpreted as a dogmatic prior where the parameters of the beta distribution C and D are chosen to deliver a constant γ. We explored a version of the model where we let the prior of γ_0 to follow Beta distribution and use Metropolis-Hastings to draw its posterior distribution. Such a specification leads to qualitative similar results but requires much longer computation time. We therefore fix the γ.

Table 5: Comparison of Individual-Level Model Fit

Moments	Horizon	Data	Benchmark	Overconfidence ($\phi = 2.25$)	Diagnostic Exp. ($\theta = 0.8$)	No LR.	No Disp.
		(1)	(2)	(3)	(4)	(5)	(6)
Panel A: CG Regression Moments							
CG a	**1-Yr**	3.828	3.080	3.201	3.086	2.063	2.956
		(0.071)	(0.061)	(0.064)	(0.061)	(0.068)	(0.058)
	2-Yr	5.032	4.558	4.581	4.546	4.145	4.569
		(0.064)	(0.056)	(0.057)	(0.056)	(0.066)	(0.052)
	3-Yr	5.253	4.991	5.003	4.986	4.827	4.947
		(0.064)	(0.058)	(0.059)	(0.059)	(0.061)	(0.053)
	4-Yr	5.063	5.008	5.017	5.003	4.953	4.941
		(0.078)	(0.073)	(0.074)	(0.074)	(0.069)	(0.066)
CG b	**1-Yr**	-0.219	0.086	-0.443	-0.312	-0.384	1.786
		(0.043)	(0.042)	(0.024)	(0.027)	(0.018)	(0.094)
	2-Yr	-0.514	-0.243	-0.413	-0.359	-0.555	0.033
		(0.046)	(0.079)	(0.060)	(0.064)	(0.022)	(0.332)
	3-Yr	-0.574	-0.365	-0.398	-0.421	-0.598	1.185
		(0.051)	(0.092)	(0.085)	(0.088)	(0.025)	(0.708)
	4-Yr	-0.418	-0.345	-0.365	-0.336	-0.601	-0.303
		(0.063)	(0.116)	(0.110)	(0.111)	(0.033)	(0.997)
Weighted Cost Function			**205.0**	**137.8**	**121.5**	**128.5**	**975.4**
Panel B: Sensitivity to Past House Price Growth							
Moments	**Horizon**	**Data**	**Benchmark**	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
Sensitivity to Past 1Q Growth	**1-Yr**	0.186	0.261	0.255	0.277	0.779	0.221
		(0.006)	(0.005)	(0.007)	(0.006)	(0.007)	(0.001)
	2-Yr	0.056	0.088	0.084	0.095	0.643	0.059
		(0.005)	(0.006)	(0.007)	(0.006)	(0.007)	(0.000)
	3-Yr	0.019	0.038	0.036	0.040	0.527	0.017
		(0.004)	(0.006)	(0.007)	(0.007)	(0.007)	(0.000)
	4-Yr	0.003	0.025	0.024	0.027	0.450	0.005
		(0.004)	(0.007)	(0.007)	(0.007)	(0.006)	(0.000)
Sensitivity to Past 2Y Growth	**1-Yr**	0.127	0.236	0.231	0.252	0.668	0.195
		(0.006)	(0.005)	(0.007)	(0.007)	(0.010)	(0.002)
	2-Yr	0.026	0.071	0.068	0.077	0.462	0.042
		(0.005)	(0.005)	(0.006)	(0.005)	(0.008)	(0.001)
	3-Yr	0.006	0.036	0.034	0.039	0.376	0.012
		(0.004)	(0.006)	(0.006)	(0.006)	(0.007)	(0.000)
	4-Yr	-0.003	0.025	0.023	0.027	0.330	0.004
		(0.004)	(0.006)	(0.006)	(0.006)	(0.007)	(0.000)
Sensitivity to Past 4Y Growth	**1-Yr**	0.131	0.231	0.226	0.248	0.632	0.181
		(0.007)	(0.006)	(0.007)	(0.007)	(0.011)	(0.003)
	2-Yr	0.019	0.077	0.074	0.083	0.421	0.039
		(0.005)	(0.005)	(0.006)	(0.005)	(0.009)	(0.001)
	3-Yr	-0.003	0.044	0.042	0.046	0.342	0.012
		(0.005)	(0.006)	(0.006)	(0.006)	(0.008)	(0.000)
	4-Yr	-0.009	0.034	0.032	0.036	0.315	0.005
		(0.004)	(0.006)	(0.007)	(0.006)	(0.007)	(0.000)
Weighted Cost Function			**116.3**	**109.4**	**131.3**	**1112.7**	**42.8**

Table 5: Comparison of Individual-Level Model Fit (cont.)

Panel C: Biases and MZ Regression Moments							
Moments	**Horizon**	**Data**	**Benchmark**	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
Forecast Biases (%)	**1-Yr**	3.593	2.876	2.889	2.793	1.803	3.007
		(0.068)	(0.063)	(0.069)	(0.065)	(0.071)	(0.062)
	2-Yr	4.935	4.450	4.466	4.429	3.979	4.474
		(0.060)	(0.056)	(0.057)	(0.056)	(0.069)	(0.051)
	3-Yr	5.240	4.994	5.006	4.989	4.838	4.961
		(0.060)	(0.057)	(0.058)	(0.058)	(0.064)	(0.051)
	4-Yr	5.072	5.011	5.020	5.004	4.990	4.942
		(0.064)	(0.064)	(0.065)	(0.064)	(0.063)	(0.057)
Inverse MZ *a*	**1-Yr**	2.234	2.513	2.528	2.362	1.147	2.786
		(0.063)	(0.053)	(0.070)	(0.064)	(0.111)	(0.029)
	2-Yr	3.227	2.956	2.964	2.953	2.459	3.236
		(0.077)	(0.069)	(0.075)	(0.071)	(0.136)	(0.011)
	3-Yr	3.511	2.863	2.876	2.897	1.947	3.233
		(0.075)	(0.081)	(0.086)	(0.082)	(0.130)	(0.004)
	4-Yr	3.702	2.905	2.905	2.907	1.320	3.253
		(0.071)	(0.084)	(0.088)	(0.085)	(0.116)	(0.002)
Inverse MZ *b*	**1-Yr**	0.124	0.203	0.199	0.237	0.564	0.143
		(0.008)	(0.006)	(0.008)	(0.008)	(0.013)	(0.003)
	2-Yr	-0.043	0.053	0.050	0.056	0.177	0.014
		(0.009)	(0.008)	(0.009)	(0.008)	(0.016)	(0.001)
	3-Yr	-0.056	0.048	0.045	0.045	0.178	0.007
		(0.008)	(0.009)	(0.010)	(0.009)	(0.015)	(0.000)
	4-Yr	-0.064	0.038	0.037	0.038	0.233	0.004
		(0.008)	(0.009)	(0.010)	(0.009)	(0.013)	(0.000)
Weighted Cost Function			**97.4**	**95.2**	**100.8**	**278.0**	**66.7**
Panel D: Forecaster Dispersion Moments							
Moments	**Horizon**	**Data**	**Benchmark**	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
Forecaster Dispersion (%)	**1-Yr**	6.098	2.841	6.282	4.669	6.065	−−
	2-Yr	4.834	3.122	3.726	3.283	4.169	−−
	3-Yr	3.519	3.507	3.941	3.579	3.345	−−
	4-Yr	2.970	3.711	4.135	3.795	2.963	−−
Weighted Cost Function			**714.7**	**359.7**	**482.7**	**109.9**	−−
Panel E: Total Weighted Cost Function							
			Benchmark	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
Total Weighted Cost Function			1133.3	702.1	836.3	1629.0	−−
Coefficient (weighted)			418.7	342.4	353.6	1519.2	1084.9
Dispersion (weighted)			714.7	359.7	482.7	109.9	−−

Notes: In Panel E, the "Total Weighted Cost Function" are calculated as follows: for the coefficients ("Coefficient") $\sum \frac{|\text{Model Implied Coefficients} - \text{Actual Coefficients}|}{\text{Actual Coefficient Standard Errors}}$,and for dispersion $\sum |\text{Model-implied Dispersion} - \text{Actual Dispersion}| \times \omega$, where the weight $\omega = \sum_{i=1}^{n} \frac{1}{\text{actual coefficient standard errors}} / n$. In Panel A through Panel C, the coefficient standard errors are in parentheses.

The main result of the paper is that the rational learning model goes a long way in explaining the key empirical moments we documented. This includes many moments commonly interpreted as providing *prima facie* evidence of deviations from rationality, such as non-trivial forecast bias, non-zero slope estimates in the CG regressions, slope coefficients in the (inverse) MZ regressions different from one, and non-trivial positive slope coefficients of forecasts on lagged realizations.

Importantly, Panel A shows that the model generates a negative slope of the individual CG regressions at horizons of 2, 3, and 4 years. The negative slope is usually interpreted as evidence that forecasters "over-react" to news that leads them to update their forecast. We discuss an alternative interpretation of this result in the context of our learning model below. We note that the CG slope is too small in absolute value, compared to the data, especially at the one year horizons. The one-year CG slope will be a moment that can be matched better once behavioral frictions are introduced.

Panel B shows that forecast sensitivity to lagged realized HPA in the model is positive, declining in the forecast horizon, and declining in the look-back window over which the regressor is computed. All of these are also features of the data. The long-run forecasts become essentially insensitive to past house price growth in both model and data. Hence the model produces the fact that forecasters are imputing "too much" mean-reversion into HPA.

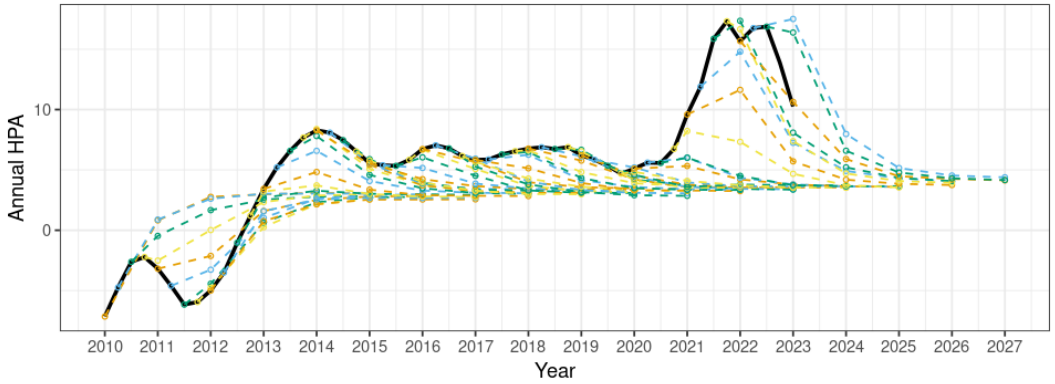
Panel C shows that our learning model produces large average in-sample forecast biases that grow with the horizon, consistent with the data. It also produces forecasts whose sensitivity to contemporaneous realized HPA (inverse MZ slope) is close to zero (and far way from one), consistent with the data .

Panel D shows that the model does a reasonably good job matching the level of forecast dispersion, especially at longer horizons of 2–4 years. The benchmark model understates the forecast dispersion at short horizons. Much more on this below.

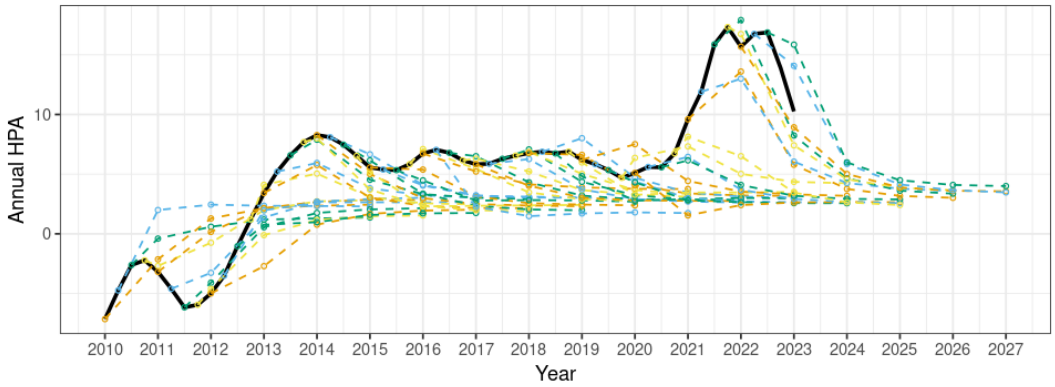
Panel E shows that the distance between model and data, as captured by the cost function, attains a value of 1,133, a reference point for other model versions discussed below.

Figure 4 plots the estimated model-implied consensus forecast for house price growth in the time series in panel (a). It looks similar to the figure in the introduction, which was based on empirical forecasts, since the model fits the data well. Panel (b) shows the same graph but for an individual forecaster, one of the 100 forecasters we have in our model simulation and estimation. The individual forecasts are naturally more volatile than the consensus. Panel (c) shows all 100 forecasters' forecasts in the model in grey, along with the consensus forecast in blue. The width of the grey area illustrate that the model is capable of generating substantial dispersion in forecasts.

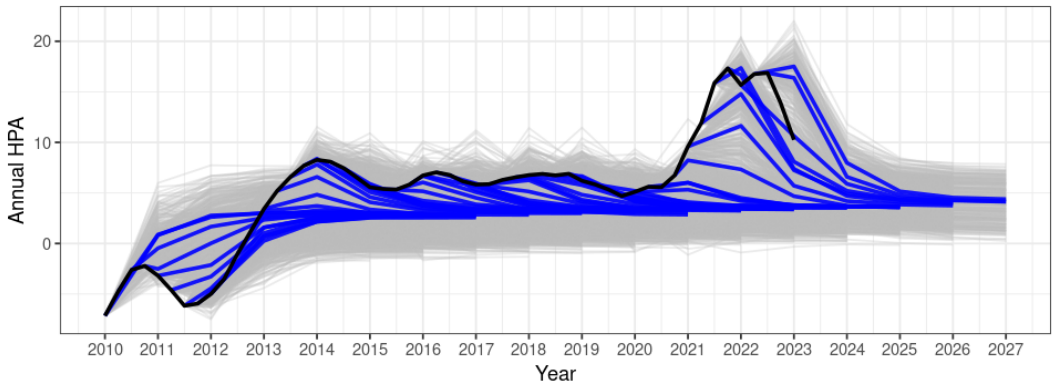
Figure 4: Model Implied Forecasts: Consensus & Individuals



(a) Model-Implied Consensus Forecast



(b) Model-implied Individual Forecast



(c) Model Implied 100 Individual Forecasts

Notes: Benchmark Rational Learning Model

Table A1 reports 24 residual standard error moments that are not included in the estimation. They measure the volatility of the regression residuals, generated both from time-series and cross-sectional volatility. They provide additional opportunities for model validation. The model fits these residual standard error

moments quite well, confirming that the model not only matches the mean of empirical moments well, but also the variation of these moments, both over time and across forecasters.

4.3 Individual versus Consensus CG Regression Slopes

Interestingly, the benchmark model also fits well the same moments obtained from consensus forecasts instead of individual forecasts (in both model and data). Table 6 shows the fit for the CG regression moments.[10] The moments in this table were not used in the estimation. The CG regression slopes are, if anything, positive when estimated on consensus data. We note the large standard errors on the consensus CG slop estimates in both model and data.

Table 6: Comparison of Consensus-Level Model Fit of CG Moments

Moments	**Horizon**	**Data**	**Benchmark**	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
		(1)	(2)	(3)	(4)	(5)	(6)
CG *a*	**1-Yr**	3.652	2.643	2.634	2.600	1.857	*2.953*
		(0.622)	(0.539)	(0.542)	(0.523)	(0.634)	(0.597)
	2-Yr	5.017	4.501	4.519	4.480	4.205	*4.565*
		(0.557)	(0.517)	(0.519)	(0.521)	(0.645)	(0.530)
	3-Yr	5.160	4.985	5.001	5.023	4.926	*4.920*
		(0.602)	(0.522)	(0.523)	(0.529)	(0.573)	(0.541)
	4-Yr	5.275	4.975	4.975	4.729	5.002	*4.932*
		(0.793)	(0.664)	(0.664)	(0.670)	(0.658)	(0.692)
CG *b*	**1-Yr**	1.808	2.287	2.344	1.770	0.239	*1.809*
		(0.971)	(0.781)	(0.797)	(0.627)	(0.388)	(0.965)
	2-Yr	-0.563	0.942	0.910	0.699	-0.774	*0.157*
		(2.150)	(2.516)	(2.613)	(2.223)	(0.557)	(3.770)
	3-Yr	-0.997	0.068	-0.285	-3.041	-1.445	*6.996*
		(3.369)	(5.830)	(5.973)	(7.634)	(0.775)	(15.285)
	4-Yr	3.571	2.890	3.618	15.753	-1.346	*3.952*
		(5.745)	(8.200)	(8.264)	(12.060)	(0.956)	(50.111)
Weighted Cost Function			**19.5**	**17.7**	**37.1**	**32.1**	**30.5**

Notes: "Weighted Cost Function" for the coefficients is $\sum \frac{|\text{Model Implied Coefficients} - \text{Actual Coefficients}|}{\text{Actual Coefficient Standard Errors}}$.

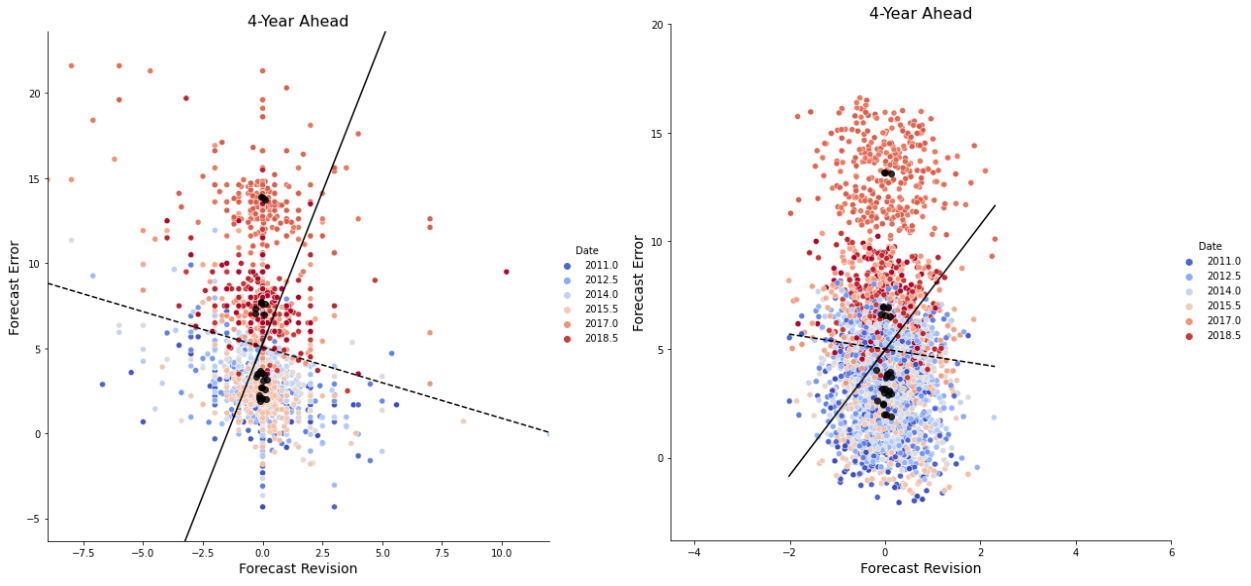
The sign switch between b^{cg} at the individual level (negative) and at the consensus level (positive) is found for many other macro-economic and financial survey data and is the topic of the paper by Bordalo et al. (2020). They argue that the overreaction to news at the individual level and the underreaction to news at the consensus level can be explained in a model where forecasters receive heterogeneous signals and have diagnostic expectations. The positive sign at the consensus level comes from aggregating forecast revisions among forecasters with heterogeneous signals, which shrinks the average revisions towards zero. The lack of reaction to other forecasters' signals creates rigidity in the consensus belief and a positive consensus coefficient. This feature arises naturally in Bayesian settings with signal dispersion, including ours.

[10] These are the only empirical moments that differ when estimated on individual versus consensus data.

They argue that the negative sign at the individual level is harder to generate, and it arises from diagnostic expectations in their setting. Our paper shows that diagnostic expectations are not necessary to generate the negative b^{cg} at the individual level. In our rational learning model, the negative slope in the individual CG regressions is not an actual overreaction to news. Agents are learning about the long-run mean of HPA in our setting, unlike in Bordalo et al. (2020). When agents receive positive news that leads them to revise their forecast upward, they optimally move their forecast up so much that it generates a positive forecast error. Given their deep uncertainty about the prior mean (high σ_μ), this "over-weighting" of the signal is optimal. Hence, learning about the long-run mean HPA in the presence of substantial parameter uncertainty generates a negative relationship between forecast error and forecast revision. More on this in the next section.

Figure 5 shows scatter plots of forecast revisions (x-axis) against forecast errors (y-axis) in the actual data (left panel) and in the benchmark model (right panel). The colored dots represent individual forecasters (organized by year) and the black dots represent the consensus (one dot per period). The dashed line shows the negative relationship based on the individual data points, while the solid black line shows the positive (and much more imprecisely estimated) relationship based on the consensus forecast.

Figure 5: Individual versus Consensus CG Regression Slopes of 4-Year Ahead Forecasts



(a) Actual

(b) Benchmark Model

Notes: Black dots show the consensus forecasts at different periods, dashed lines show the individual CG regressions, and solid lines show the consensus CG regressions.

4.4 Understanding the Main Economic Forces

Two key economic forces drive our model's forecasts at the individual level and across forecasters. We discuss these economic forces and demonstrate their presence in our model through simulations.

4.4.1 Learning About the Long-Run Mean and Individual Forecasts

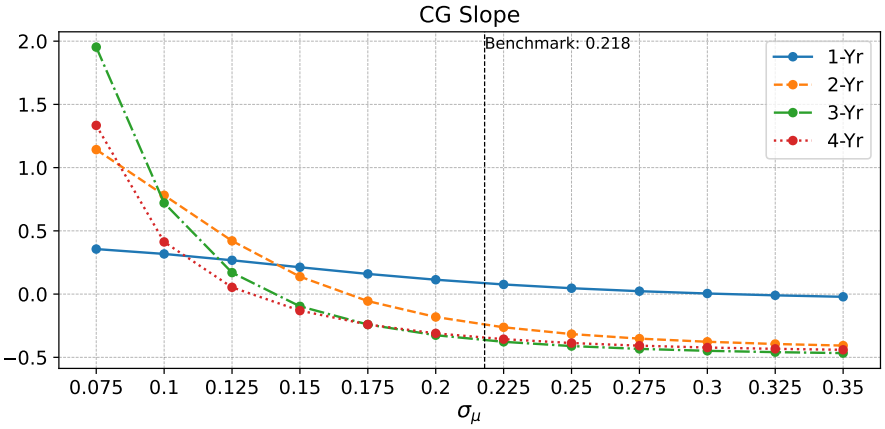
Individual forecasters in our model do not know the true value of long-run house price growth ($\overline{\mu}$), and must learn about it by combining prior information and signals using Bayes' Law. This learning about long-run mean growth generates patterns that appear anomalous from the perspective of a rational expectations model. The magnitude of σ_μ, which governs the precision of forecasters' prior beliefs, determines the volatility of their forecasts. This volatility in turn prominently affects the CG coefficients and the sensitivity of forecasts to past realized growth.

If forecasters have more diffuse priors (larger σ_μ), they assign more weight to their signals when making forecasts. In this case, this "modest" forecaster's forecasts may appear excessively volatile or overreacting to news about house price growth. An econometrician would observe more negative individual CG coefficients and a stronger response of the forecast to past growth. Conversely, more "dogmatic" forecasters with tighter prior beliefs (smaller σ_μ) assign less weight to data in their forecasts. Consequently, their forecast revisions do not predict future forecast errors, and their response to past realized growth appears too sluggish compared to the observed auto-correlation in realized house price growth. In the extreme case where forecasters think they know the exact underlying value of the long-run house price growth ($\sigma_\mu = 0$), they do not revise their long-run forecasts at all, and the sensitivity of long-run forecasts to past returns is zero. When their beliefs are consistent with the true parameter values, we are back to the Full Information Rational Expectation (FIRE) case.

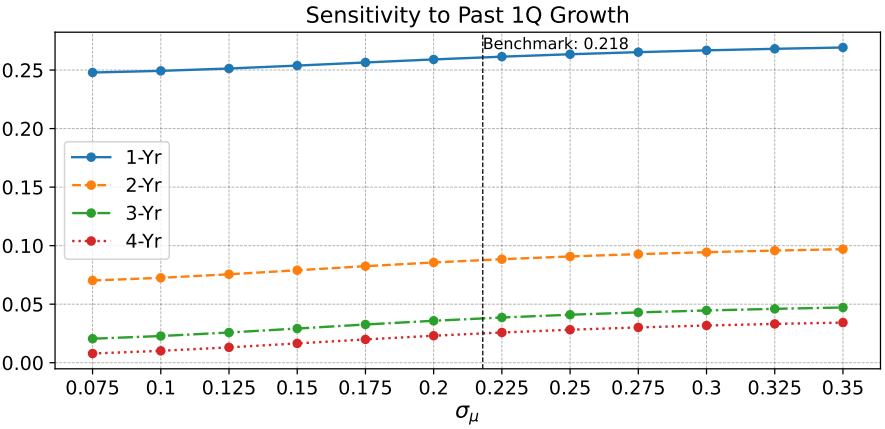
Figure 6 demonstrates this mechanism based on model simulations for different values of σ_μ, reported on the x-axis. On the y-axes, it plots the values of the individual CG coefficients (Panel a), the sensitivity of forecast to past 1Q growth (Panel b) and to past 4Y growth (Panel c). As the value of σ_μ becomes smaller than the benchmark value, the CG coefficients gradually increase, and become positive around $\sigma_\mu = 0.15$, mimicking the forecasts of a "dogmatic" forecaster. Conversely, the individual CG regression slope becomes more negative as σ_μ becomes larger, i.e., as the forecaster becomes more "modest". The effects on forecasters' sensitivity to past realized growth are the opposite. As σ_μ increases, forecast sensitivity to past realized growth increases. Notably, the effect of σ_μ is the largest on the long-run forecasts and long-run CG slopes, confirming our "learning about the *long-run* mean" mechanism.

Figure 6 also reveals how the two sets of empirical moments, individual CG coefficients and the sensitiv-

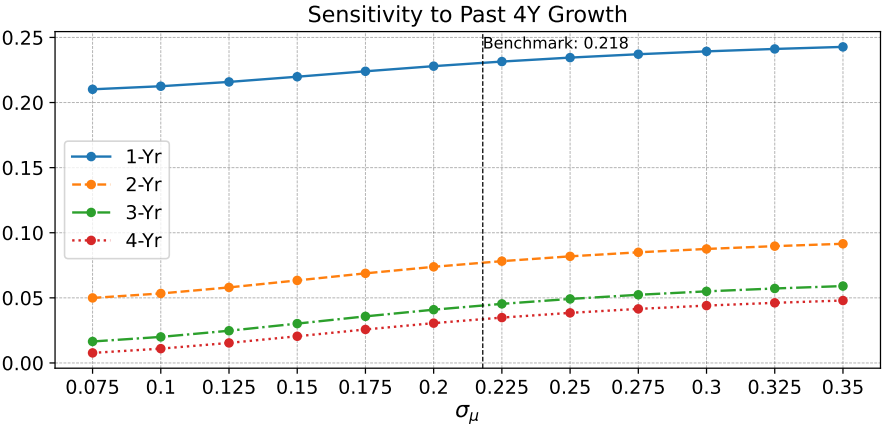
Figure 6: Impact of Forecaster Conviction (σ_μ) on Individual Forecasts



(a) Impact of σ_μ on CG-b coefficients



(b) Impact of σ_μ on sensitivity to 1Q past growth



(c) Impact of σ_μ on sensitivity to 4Y past growth

Notes: The graph varies the parameter σ_μ, holding all other parameters fixed at the benchmark model's values.

ity to past growth, help identify the parameter σ_μ in our model.[11] Forecasters in the data exhibit negative CG coefficients, and a substantial positive short run but zero long-run response to past realized growth. These moments point towards an intermediate value for σ_μ. A higher value than 0.25 would help to bring the model closer to the data in terms of the individual CG slopes at horizons 2-4, but would hurt the model's fit for the sensitivity of 2–4-year-ahead forecasts to lagged HPA. The resulting intermediate value of 0.25 implies a high degree of prior uncertainty about long-run HPA. This is a world in which learning about long-run growth matters (far from the dogmatic case), and new data can really sway forecasts.

Our learning environment is similar to Collin-Dufresne, Johannes and Lochstoer (2016), where learning about the long-run mean in a multi-dimensional setting leads to "excessive" volatility in the subjective expectations of forecasters. Their model does not consider the implications for the CG regression nor the implications for the term structure of sensitivities to past realized house price growth.

4.4.2 Uncertain vs. Biased Initial Beliefs

The economic mechanism behind our model is different from the one proposed in Farmer, Nakamura and Steinsson (2023). In Farmer, Nakamura and Steinsson (2023), forecasters have biased initial beliefs. Specifically, when their prior beliefs about the persistence parameter $\overline{\rho}$ are upward-biased, their model can generate "anomalous" empirical moments, including the negative individual CG slopes.

Our model does not require forecasters' beliefs to be biased. Instead, the forecasters in our model are merely *uncertain* about the long-run mean of house price growth. Even though biased beliefs and uncertain beliefs can both generate negative CG slopes, only uncertain beliefs can simultaneously match the forecasts' sensitivity to past realized house price growth (as explained in the previous subsection).

To highlight the role of biased beliefs about persistence, we vary the prior mean μ_ρ from 0.355 to 0.95, holding all other parameters at their benchmark estimates. Panel (a) of Figure 7 shows the individual CG slope while panels (b) and (c) show the sensitivity of the forecast to past realized house price growth.

Higher values of μ_ρ result in more negative individual CG slopes. This is consistent with the findings of Farmer, Nakamura and Steinsson (2023).[12] The decrease in CG slopes is most prominent at the 1-year horizon. Indeed, when μ_ρ is around 0.85, the CG slopes are closer to the data than at the benchmark estimate of $\mu_\rho = 0.66$, even though the overall model fit deteriorates significantly.

An important reason for the worse fit arises from the model's prediction for the forecast sensitivity to past realized house price growth. The bottom two panels demonstrate that these sensitivities increase in

[11] Appendix E.1 provides a more formal discussion of parameter identification. Inspired by Andrews, Gentzkow and Shapiro (2017), it varies each parameter by 20% up and down around its point estimate, and shows his this variation affects each of the moments used in estimation.

[12] Higher values for the prior mean of μ_ρ, such as values above 0.9, also result in higher values for the posterior mean, and a higher upward-biased estimate assuming the true coefficient $\overline{\rho}$ takes on some intermediate value around 0.7.

the value of μ_ρ. As μ_ρ increases to, say, 0.85, forecasters become much more sensitive to past realized price growth, not only relative to the benchmark model but also relative to the data. If forecasters' initial beliefs about μ_ρ are too high, forecasters consider recent house price growth to be more persistent than it actually is. Their forecasts become too sensitive to past realized growth.

Farmer, Nakamura and Steinsson (2023) do not consider the forecasts' response to past realized growth as part of their empirical moments (and they do not study house price data in the first place). Once this important set of moments is included–at least in the housing context–the data favor lower values for μ_ρ, as shown in panel (d).

4.4.3 Prior and Signal Heterogeneity and Forecast Dispersion

The second set of key targets for our model is the cross-sectional dispersion in forecasts. Both heterogeneity in forecasters' prior beliefs (governed by λ) and heterogeneity in their signals (governed by Γ) generate forecast dispersion. Our model reveals how these two different forces affect both the level and the term structure of the forecast dispersion, adding new insight to the previous literature (e.g., Patton and Timmermann, 2010).

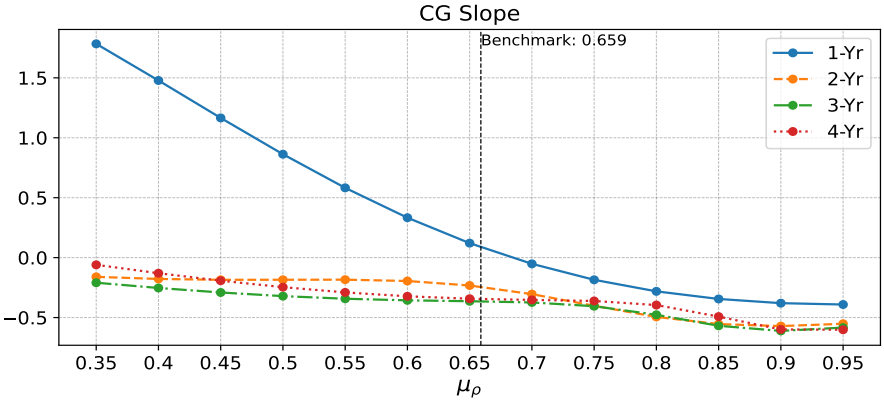
First, the impact of signal heterogeneity ($\overline{\Gamma}$) on the level of forecast dispersion is limited in a rational learning model. When forecasters are rational and aware that their signals contain a lot of noise (high $\Gamma = \overline{\Gamma}$), they under-weight the signal when forming expectations following Bayes' Rule. An increase in Γ initially increases forecast dispersion due to learning from data. However, its impact diminishes as $\overline{\Gamma}$ increases further as forecasters discount their own signal when they realize that $\overline{\Gamma}$ is high.

Panel (a) of Figure 8 confirms this mechanism. It plots forecast dispersion, the cross-sectional variance of the forecast, for different values of the parameter $\overline{\Gamma}$. The declining marginal impact of higher signal dispersion is visible at all horizons, but is especially strong for the 1-year forecast dispersion which displays a hump shape. At low levels of $\overline{\Gamma}$, increases in $\overline{\Gamma}$ lower the slope of the term structure of forecast dispersion, while beyond the benchmark value of $\overline{\Gamma} = 1.25$, increases in $\overline{\Gamma}$ result in a steeper term structure of forecast dispersion.

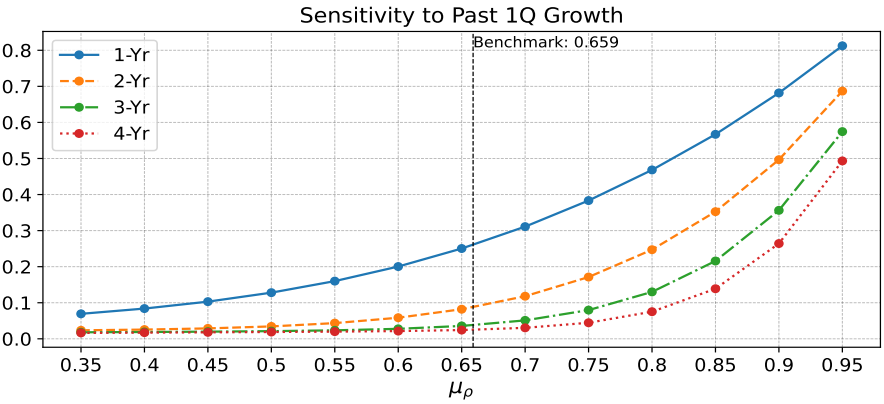
Second, the heterogeneity in forecasters' prior beliefs about the long-run mean (λ) has a first-order impact on both the level and the slope of the term structure of forecaster dispersion. A higher value of λ increases the level of forecast dispersion at all horizons. It does so more for the long-run forecasts, leading to an upward-sloping term structure of forecast dispersion. Panel (b) of Figure 8 confirms this effect given our benchmark estimate of $\sigma_\mu = 0.218$.

Finally, the effect of λ on forecast dispersion should be jointly considered with the level of forecaster

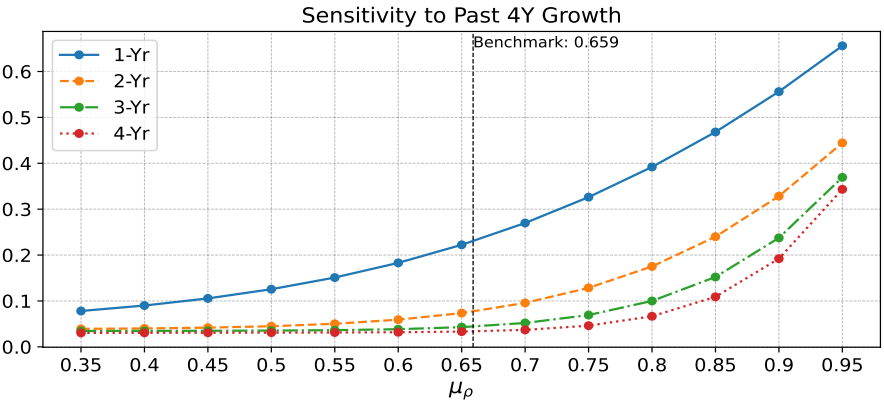
Figure 7: Impact of Prior Belief of the Persistence Parameter (μ_ρ) on Individual Forecasts



(a) Impact of μ_ρ on CG-b coefficients



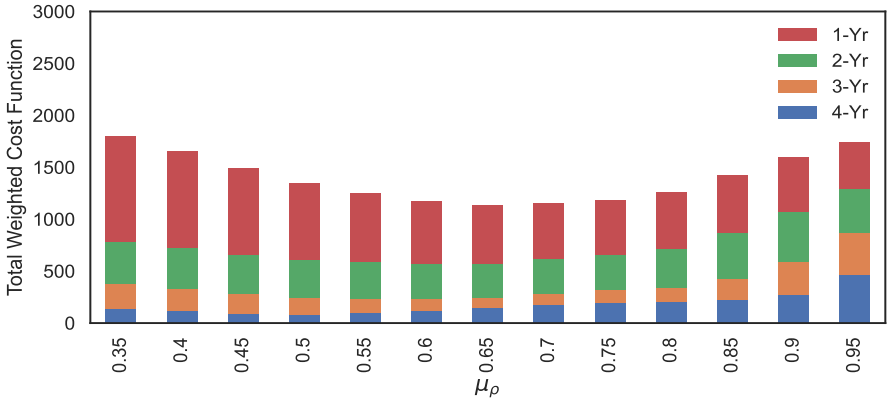
(b) Impact of μ_ρ on sensitivity to 1Q past growth



(c) Impact of μ_ρ on sensitivity to 4Y past growth

conviction of their prior beliefs (σ_μ). For more "modest" forecasters who place a lower weight on their prior beliefs in their expectation formation process (high σ_μ), the impact of λ on the level of dispersion is limited. Panel (c) of Figure 8 shows that, when we increase σ_μ to 0.5, the increase in dispersion is much smaller

Figure 7: Impact of Prior Belief of the Persistence Parameter (μ_ρ) on Individual Forecasts (cont.)



(d) Impact of μ_ρ on total weighted cost function

Notes: The graph varies the parameter μ_ρ, holding all other parameters fixed at the benchmark model's values.

whenλ increases. We also notice a much lower level of forecast dispersion at all horizons. These results help explain the intermediate value for σ_μ found by the estimation routine.

In sum, generating enough forecast dispersion requires enough dispersion in the prior (a high value of λ) as well as a modest amount of uncertainty about the long-run mean (a low enough value of σ_μ), as well as enough signal dispersion $\overline{\Gamma}$. There are limits to how much forecast dispersion can be generated by cranking up signal dispersion when investors are rational (know the true signal noise).

5 The Impact of Behavioral Biases

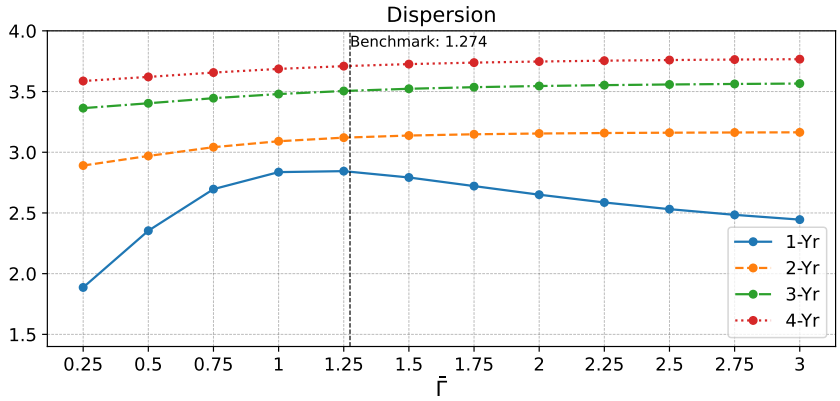
We now enrich the learning model with one of two different behavioral frictions and show that these frictions help to further improve the fit along two important dimensions: the short-run CG regressions slope and the short-run cross-sectional forecast dispersion.

5.1 Overconfidence

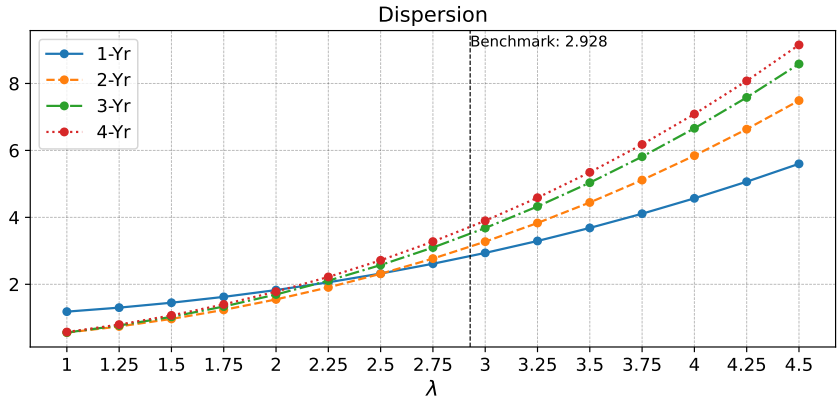
We allow for a difference in the true and the perceived signal variance, modulated by the parameter ϕ. The rational learning benchmark, where the two are the same, is $\phi = 1$. When $\phi > 1$, the true signal variance is higher than the perceived one. In this case, forecasters are over-confident about the informational content of the signal. The opposite is true for $\phi < 1$.

Panel (a) of Figure 9 shows the model-implied CG slope as we change the value of ϕ. As ϕ increases from 0.25 (under-confidence) to 3 (over-confidence), the CG slope based on 1-year forecasts changes dramatically

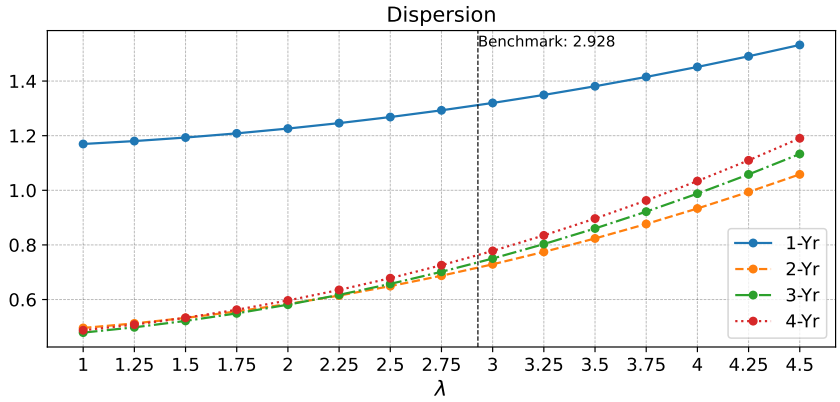
Figure 8: Impact of Signal Dispersion ($\overline{\Gamma}$) and Prior Heterogeneity (λ) on Forecast Dispersion



(a) Impact of $\overline{\Gamma}$ on forecast dispersion



(b) Impact of λ on forecast dispersion when $\sigma_\mu = 0.218$



(c) Impact of λ on forecast dispersion when $\sigma_\mu = 0.5$

Notes: Panel (a) varies the signal variance parameter $\Gamma = \overline{\Gamma}$, holding all other parameters fixed at the benchmark model's values. Panel (b) varies the parameter λ, which governs the dispersion in the prior about the long-run mean house price growth rate, holding all other parameters fixed at the benchmark model's values. Panel (c) redoes panel (b) for a value of $\sigma_\mu = 0.5$, higher than the benchmark value of 0.218.

from +0.9 to -0.5. On the other hand, the CG slopes based on forecasts of longer horizons do not change much at all. Forecasts of 3-4 year still exhibit negative CG slopes, due solely to forecasters learning about the long-run mean. This result is intuitive. A large ϕ should lead to a more negative CG regression slope as over-confident forecasters revise their forecasts *too much* compared to the true information content of their signals. Furthermore, the effect of ϕ should be most visible in the short-term forecasts, as private signals mostly concern the current period. After all, everybody agrees on the true, past house price growth, limiting the potency of over-confidence to affect longer-run forecasts.

Panel (b) of Figure 9 demonstrates how the cross-sectional forecast variance changes with ϕ. As ϕ increases, forecast dispersion of all horizon increases. However, the 1-year forecast dispersion displays by far the strongest impact . Intuitively, overconfidence should increase the short-term forecast dispersion. As we discussed in Section 4.4.3, signal heterogeneity has limited impact on forecast dispersion because rational forecasters discount their signals when forming forecasts if their signals contain a lot of noise. However, when forecasters over-estimate the precision of their signals, they discount the signal less than they should. In this setting, higher signal heterogeneity ($\overline{\Gamma}$) *does* translate into higher forecast dispersion.

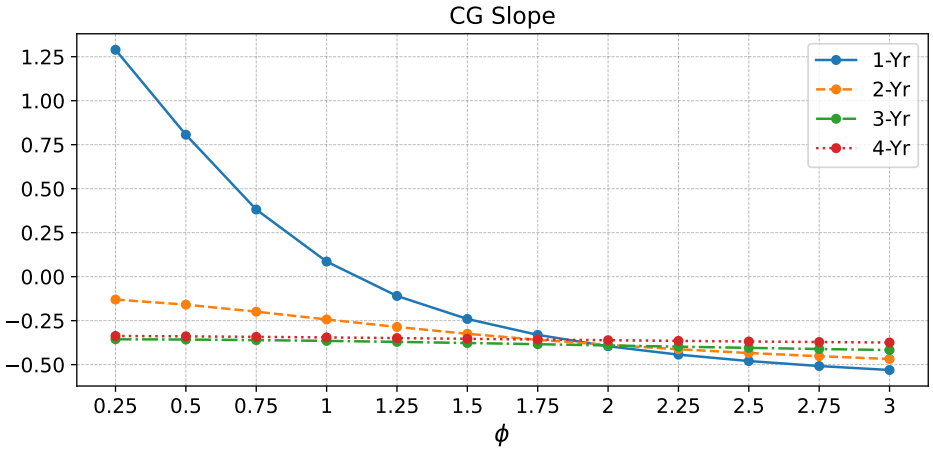
5.2 Diagnostic Expectations

We allow forecasters to display diagnostic expectations, modulated by the behavioral filter parameter θ as proposed in Bordalo et al. (2020). Figure 10 shows how the individual CG regression slope (Panel a) and the cross-sectional forecast variance (Panel b) are affected by changing θ up or down by 0.2 increments. The rational learning benchmark model sets $\theta = 0$.

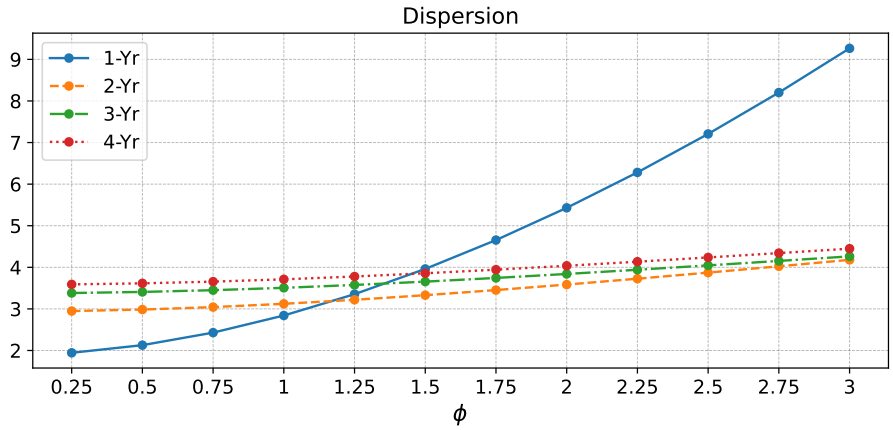
The diagnostic filter affects short forecasting horizons (1-to-2 years) more than long horizons (3-4 years). This is intuitive given that the behavioral filter affects the dynamics of the latent state and the latent state has a modest half-life.

While the rational learning model can generate a negative slope coefficient in the 1-year CG regression, at least for some combinations of parameters (different from the benchmark model), the distance between the one-year CG b coefficient in the benchmark model and the data is substantial. Consistent with the intuition in Bordalo et al. (2020), diagnostic expectations can help generate a *strongly* negative CG short-horizon slope coefficient. In fact, the best fit for the one-year CG slope and cross-sectional forecast variance are obtained for a value around $\theta = 0.8$, close to the values proposed by Bordalo et al. (2020). Diagnostic expectations help resolve the tension between short-run and long-run forecast dispersion. Learning in the presence of uncertainty about the long-run mean growth rate $\overline{\mu}$ generates high long-run forecast dispersion, while diagnostic expectations helps generate the high short-run forecast dispersion.

Figure 9: The Impact of Overconfidence (ϕ)



(a) Impact of ϕ on CG-b coefficients



(b) Impact of ϕ on forecast dispersion

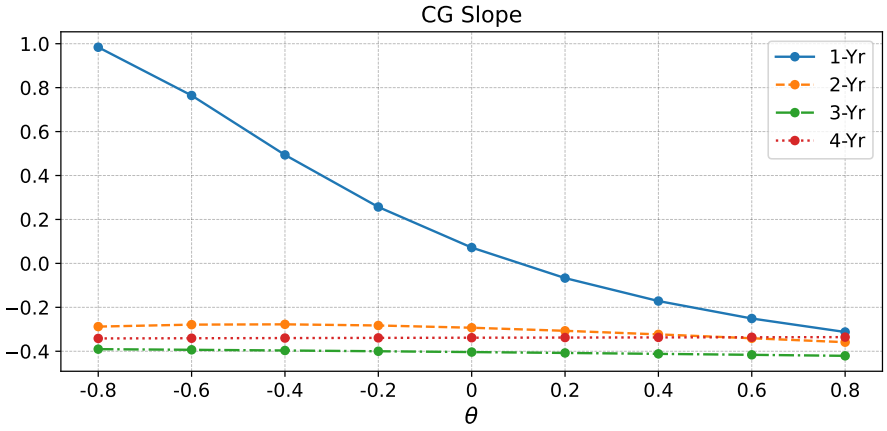
Notes: The model with overconfidence introduces a wedge between the objective signal variance $\overline{\Gamma}$ and the subjective variance Γ. The parameter ϕ controls the ratio of the former to the latter. All other parameters are kept at their benchmark values.

5.3 The Impact of Behavioral Biases on Overall Model Fit

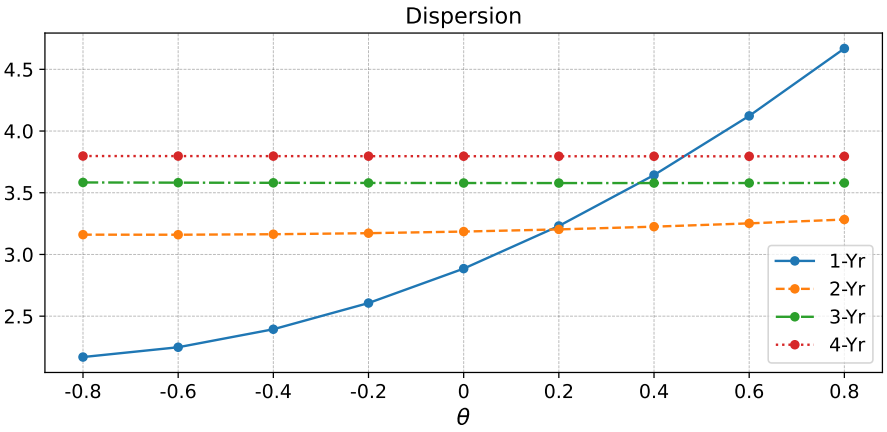
Figure 11 compares the cost function values for different levels of overconfidence (panel a) or diagnostic expectations (panel b). Incorporating either overconfidence ($\phi > 1$) or a diagnostic filter ($\theta > 0$) improves the model fit, i.e., reduces the total cost function. Specifically, an overconfidence level at $\phi = 2.25$ and a diagnostic expectation of $\theta = 0.8$ can reduce the cost function by 38% and 26%, respectively, bringing our model significantly closer to the data.

Columns 3 and 4 of Table 5 show that the reduction in total costs comes from both the reduction in regression coefficient estimates (first 32 moments), and the forecast dispersion (4 moments). The improvement comes mainly from the short-term CG slopes and short-term dispersion. Figure 11, which also de-

Figure 10: The Impact of Diagnostic Expectations (θ)



(a) Impact of θ on CG coefficients



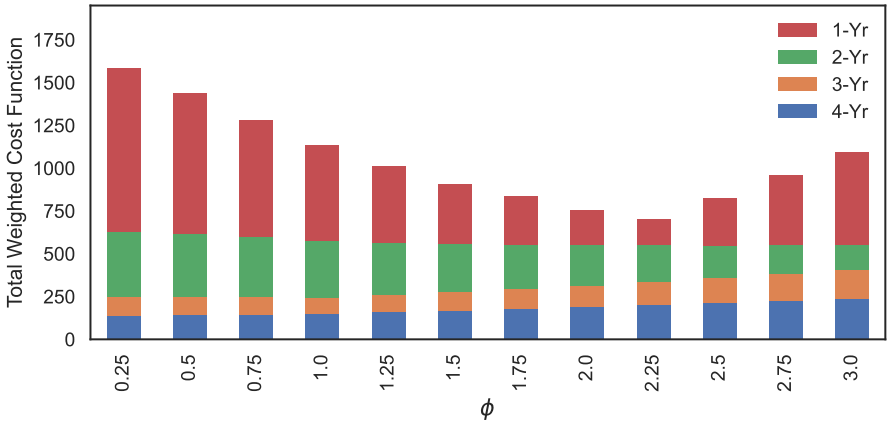
(b) Impact of θ on Dispersion

Notes: Based on the benchmark model.

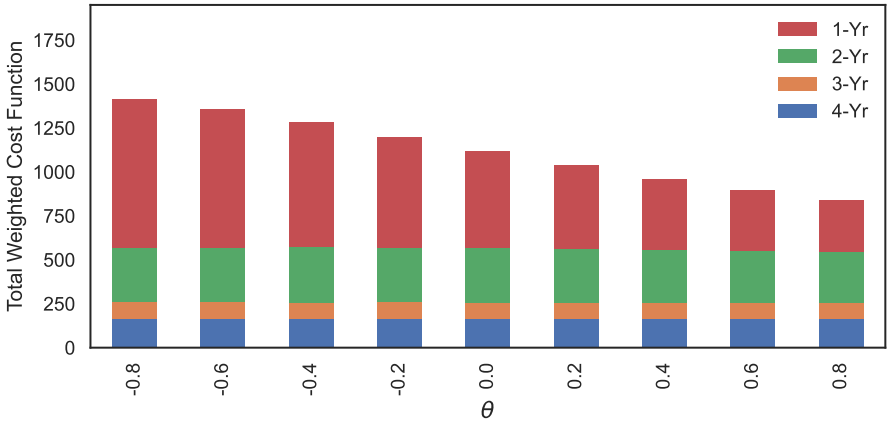
composes the total cost functions into the contributions from each forecast horizon, shows that nearly the entire improvement comes from the one-year-horizon forecasts. The long-run moments, on the other hand, are not much impacted by either behavioral frictions.

These results highlight the importance of having a rich enough rational benchmark model when making inference on the rationality of subjective forecasts. With the rational learning benchmark, we can make quantitative inference on *how much* and *where* behavioral biases play a role in shaping expectations. When added to a full-fledged learning model with uncertainty about the long-run mean, either overconfidence or diagnostic expectations are very helpful to generate realistic short-run CG coefficients and short-run forecast dispersion. These biases have little impact on the longer-run moments, which are successfully accounted for by learning dynamics.

Figure 11: The Impact of Behavioral Biases on Overall Model Fit



(a) Impact of ϕ on total weighted cost function



(b) Impact of θ on total weighted cost function

Notes: The height of the bars shows the overall cost function value, the goodness of fit of the 36 moments in Table 5. It breaks down this cost function into the contribution form each of the four different forecast horizons. Panel (a) is for various values of the overconfidence parameter ϕ, panel (b) for the diagnostic expectations parameter θ. All other parameters are held at their benchmark model's values.

6 Assessing Importance of Term Structure and Cross-sectional Data

Most standard survey data sources contain limited information on the term structure of forecasts and/or limited information on the cross-section of forecasters. Indeed, the typical survey data would contain only a one-period ahead consensus forecast. This section discusses the importance of term structure and cross-sectional information for parameter inference, and hence for the interpretation of forecast data.

6.1 Benchmark Model Estimation without Term Structure

We consider a special case of the benchmark model where we have no term structure of forecasts. We only have access to the one-year ahead forecast ($h = 1$). The estimation of this model involves 9 rather than 36 moments. Its purpose is to illustrate the role of term structure of forecaster data.

Parameter Estimates The parameter estimates are in the column "No LR." of Table 4. This model estimates a somewhat lower prior mean for house price growth (μ_μ), with somewhat less dispersion across forecasters (λ), and somewhat less uncertainty (σ_μ). However, the main difference is the lower estimate for the prior mean of the persistence parameter ($\mu_\rho = 0.54$ rather than 0.66) with a much larger prior standard deviation ($\sigma_\rho = 0.182$ rather than 0.023). Finally, this model estimates a similar amount of signal dispersion across forecasters than in the benchmark model. This shows that the parameter estimates are significantly different when term structure information is omitted or not available. Put differently, the lack of long-run forecasts can lead to severely biased inference on parameters.

Model Fit The fifth column of Table 5 shows the resulting model fit for the model without long-run forecast moments (column "No LR."). It reports all model-implied moments even though only the 1-year moments are used in estimation. The overall fit for all 36 moments is much worse than in the benchmark model; the distance between the model-implied and observed moments rises from 1,133 to 1,629. This shows that there is important information in long-horizon forecasts that does not naturally get captured when estimating the model on one-year data and studying its implications for longer forecast horizons.

This model does an excellent job fitting the forecast dispersion across horizons, and a good job fitting the individual CG moments across horizons. However, this comes at the expense of a much worse fit for the sensitivity of forecasts to lagged house price growth, which the model severely overstates.

This excess sensitivity arises from the much higher value for the posterior persistence parameter in the no LR model.[13] In turn, this high posterior estimate results from the much higher prior uncertainty about the persistence parameter (σ_ρ is 8 times higher than in the benchmark).[14] Intuitively, when agents are more uncertain about the persistence parameter, they are likely to revise it more strongly. In the benchmark model, the prior on persistence is fairly tight, limiting the role for learning about $\overline{\rho}$.

[13] Indeed, the value of the posterior persistence parameter in the benchmark model (Figure A2 panel B) is much lower than the posterior persistence parameter in the no LR model (Figure A3 panel B).

[14] The prior mean estimate (μ_ρ) is actually lower than in the benchmark, but this effect gets overpowered by the uncertainty channel. Panel (d) of Figure A4 and A5 In Appendix E.1 provide quantitative analysis to compare the impact of μ_ρ and σ_ρ on the sensitivity of forecasts to lagged house price growth.

6.2 The Impact of Forecaster Heterogeneity

We consider a second special case of the benchmark model where instead of using the individual forecaster panel data, we only use the consensus forecast data, i.e., the average across the forecasters. This case serves as a tool for explaining the role of forecaster heterogeneity. It allows us to speak to the literature that studies individual versus consensus forecasts (e.g., Bordalo et al., 2020). We use 32 moments in the estimation of this model. Compared to the benchmark model, we drop the four moments that measure forecast dispersion (since there is no dispersion in this model). These 32 moments are the same as in the benchmark model, except that they are based on consensus forecasts rather than individual forecasts. This only makes a difference for the CG slope, which–as we pointed out–changes sign when going from individual to aggregate forecasts.

Parameter Estimates Table 4 shows the parameter estimates from the simpler model estimated on consensus data in the column labeled "No Disp". The main difference with the benchmark model is the much lower estimate for prior uncertainty about the long-run mean house price growth rate, σ_μ. The latter is 0.04 in the No Disp model compared to 0.22 in the benchmark model.

Section 4 extensively discussed the crucial role of σ_μ. Without a high enough value for σ_μ, beliefs about mean house price growth become too "dogmatic." These parameter values effectively shut down the "learning about the long-run mean" channel. This channel is the key driver of the negative CG regression slopes at horizons $h = 2 - -4$.

Model Fit Table 5 shows the resulting model fit for the model without investor heterogeneity (column "No Disp."). The overall fit for the 36 moments deteriorates dramatically from 1,133 to 3,260.

Not surprisingly, the model produces no cross-sectional forecast dispersion. As anticipated, the model without forecaster dispersion cannot produce negative individual CG slopes because of its low value for σ_μ. Rather the model implies similar CG slopes at the individual and the consensus level. Since the latter are positive in the data, and well-matched (see Table 6), the No Disp model implies the wrong sign for the individual CG slopes. The remaining moments, including the sensitivity of forecasts to lagged house price growth, are matched reasonably well.

In sum, having access to panel survey data is a distinct advantage when it comes to identifying the important role of learning about the long-run mean, as revealed by the negative sign of the long-run CG regression slopes.

7 Conclusion

We argue that a rational learning model with parameter uncertainty about the long-run mean and forecaster heterogeneity goes a long way towards accounting for the features of forecast data traditionally associated with behavioral biases. We do so in the context of U.S. house price forecasts, exploiting a unique panel data set of individual-level forecasts that contain a term structure of forecasts. Theories of behavioral biases such as overconfidence and diagnostic expectations are useful complements that help the model account for the high degree of over-reaction and dispersion in short-horizon forecasts.

Future work could explore additional heterogeneity informed by observable forecaster characteristics. Estimating our model on macro-economic and financial survey data is another exciting avenue for future work.

References

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro. 2017. "Measuring the Sensitivity of Parameter Estimates to Estimation Moments." *The quarterly journal of economics*, 132(4): 1553–1592.

Armona, Luis, Andreas Fuster, and Basit Zafar. 2019. "Home price expectations and behaviour: Evidence from a randomized information experiment." *The Review of Economic Studies*, 86(4): 1371–1410.

Bailey, Michael, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel. 2019. "House price beliefs and mortgage leverage choice." *The Review of Economic Studies*, 86(6): 2403–2452.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. 2018. "The economic effects of social networks: Evidence from the housing market." *Journal of Political Economy*, 126(6): 2224–2276.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2015. "X-CAPM: An extrapolative capital asset pricing model." *Journal of financial economics*, 115(1): 1–24.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *The American Economic Review*.

Case, Karl E, and Robert J Shiller. 1988. "The behavior of home buyers in boom and post-boom markets."

Case, Karl E, Robert J Shiller, Anne K Thompson, David Laibson, and Paul Willen. 2012. "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets." *Brookings Papers on Economic Activity*, 265–315.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What can survey forecasts tell us about information rigidities?" *Journal of Political Economy*, 120(1): 116–159.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A Lochstoer. 2016. "Parameter Learning in General Equilibrium: The Asset Pricing Implications." *The American Economic Review*, 106(3): 664–698.

De Stefani, Alessia. 2021. "House price history, biased expectations, and credit cycles: The role of housing investors." *Real Estate Economics*, 49(4): 1238–1266.

Farmer, Leland, Emi Nakamura, and Jón Steinsson. 2023. "Learning About the Long Run." *The Journal of Political Economy*.

Favilukis, Jack, Sydney C Ludvigson, and Stijn Van Nieuwerburgh. 2017. "The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium." *Journal of Political Economy*, 125(1): 140–223.

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

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar. 2022. "Expectations with endogenous information acquisition: An experimental investigation." *Review of Economics and Statistics*, 104(5): 1059–1078.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021*a*. "Five facts about beliefs and portfolios." *American Economic Review*, 111(5): 1481–1522.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021*b*. "The joint dynamics of investor beliefs and trading during the COVID-19 crash." *Proceedings of the National Academy of Sciences*, 118(4): e2010316118.

Glaeser, Edward L, and Charles G Nathanson. 2017. "An extrapolative model of house price dynamics." *Journal of Financial Economics*, 126(1): 147–170.

Greenwald, Daniel L, and Adam Guren. 2021. "Do credit conditions move house prices?" National Bureau of Economic Research.

Greenwood, Robin, and Andrei Shleifer. 2014. "Expectations of returns and expected returns." *The Review of Financial Studies*, 27(3): 714–746.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante. 2020. "The housing boom and bust: Model meets evidence." *Journal of Political Economy*, 128(9): 3285–3345.

Kuchler, Theresa, Monika Piazzesi, and Johannes Stroebel. 2023. "Housing market expectations." In *Handbook of Economic Expectations*. 163–191. Elsevier.

Liu, Haoyang, and Christopher Palmer. 2021. "Are stated expectations actual beliefs? New evidence for the beliefs channel of investment demand." National Bureau of Economic Research.

Mincer, Jacob A, and Victor Zarnowitz. 1969. "The evaluation of economic forecasts." In *Economic forecasts and expectations: Analysis of forecasting behavior and performance*. 3–46. NBER.

Patton, Andrew J, and Allan Timmermann. 2010. "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion." *Journal of Monetary Economics; Amsterdam*, 57(7): 803.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of monetary Economics*, 50(3): 665–690.

Van Nieuwerburgh, Stijn, and Laura Veldkamp. 2009. "Information immobility and the home bias puzzle." *The Journal of Finance*, 64(3): 1187–1215.

Woodford, Michael. 2009. "Information-constrained state-dependent pricing." *Journal of monetary economics*, 56: S100–S124.

Appendix

A The Timing of Pusenomics Survey

The Pulsenomics Survey asks the its participants to make forecasts for the end-of-year price growth of different forecasting horizons at quarterly frequency. We therefore design the model and notation to fully take advantage of the survey structure.

First, define the log price growth, or log price returns

$$y_{T,h} = log\left(\frac{P_T}{P_{T-h}}\right)$$

where T denotes the date at which the end price is collected and h the holding period for the price return.

As a background of the Paulsenomics Survey, the survey for each quarter is conducted in the middle of the second month of each quarter. For example, the survey for 2020-Q1 is conducted in the middle of February. Zillow Home Value Index (ZHVI) for a given month is published on the third Thursday of the following month. Thus, the latest available ZHVI we have for 2020-Q1 Survey is ZHVI of Dec 2019. Therefore, in the 2020-Q1 survey, the latest information that experts actually use is from 2019-Q4. Suppose the survey is conducted in year y_s, quarter q_s, e.g., 2020Q1. Then experts actually form expectation using the data of 2019Q4,

- t are the end of quarter dates of the experts' information set, e.g. $t = 2019Q4$;
- $k = 4, 8, 12, 16, 20$ are forecasting horizons, e.g. $k = 4$ is the current year forecast (4 quarters ahead), $k = 8$ the next one year forecast and etc;
- $q_s = 1, 2, 3, 4$ are the calendar quarters in which the surveys are conducted in real time;
- $q = q_s - 1$ is the effective quarter end, of which the experts' make their forecasts based on.

For our survey, $h := 4$, since the expected price growth is always measured over a 4-quarter horizon, starting from Q1 to Q4 each calendar year. And the price growth always takes the the form of

$$y_{t+k-q,4} = log\left(\frac{P_{t+k-q}}{P_{t+k-q-4}}\right)$$

As a demonstrating example for why the quarter in which the survey is conducted matters, consider the next year (8Q ahead) forecast, conducted at 2020-Q2, so the information set for forecasters is up until 2020-Q1. This would be about the price growth from 2020-Q4 (2 quarters from now) to 2021-Q4 (6 quarters from

now). Indeed, we have

$$t + k - q : 2020\text{Q1+}8Q - 1Q = 2021Q4$$

Notice that for the current year forecast ($k = 4$), at the time that forecasts are made, there had already been realized price growth observed within the forecasting horizon, which will be taken into account in the framework.

Define the (subjective) expectation over the next quarter, made by forecaster i, as

$$s_{i,t} := E_{i,t}(y_{t+1,1}) = E_{i,t}\left[log(\frac{P_{t+1}}{P_t})\right]$$

The model implies that the m period ahead forecasts are given by

$$E_{i,t}(y_{t+m,1}) = \mu_{i,t} + \rho_{i,t}^m x_t \tag{A1}$$

so the μ_t will be the long-run forecast as $m \to \infty$, while x_t will be the transitory component that drives part of the short-term forecasts.

B Posterior Distributions of Parameters and States

B.1 Model Setup

$$\begin{aligned}
y_t &= \mu + x_t + e_t,\ e_t \sim N(0, H_t) \\
x_t &= \rho x_{t-1} + \sqrt{1-\gamma}\sigma\eta_t,\ \eta_t \sim N(0,1) \\
t &= 1, ..., T \\
H_t &= \begin{cases} \gamma\sigma^2 & , t \neq T \\ (\gamma + \Gamma^2)\sigma^2 & , t = T \end{cases}
\end{aligned}$$

The individual subscripts, i, are ignored.

B.2 Prior Beliefs of Parameters and States

Parameters:

$$\mu \sim N(\mu_\mu, \sigma_\mu^2)$$
$$\rho \sim N(\mu_\rho, \sigma_\rho^2)$$
$$\sigma^2 \sim IG(A, B)$$

States:

$$x_1 \sim N(0, 1)$$

B.3 Gibbs Sampling

We start with an initial guess of the parameter $\theta^{(0)} = (\mu^{(0)}, \rho^{(0)}, \sigma^{(0)}, x_{1:T}^{(0)})'$. Given a draw of the parameters $\theta^{(b)}$, we draw $\theta^{(b+1)}$ as follows:

1. Draw $\rho^{(b+1)}|\mu^{(b)}, \sigma^{(b)}, x_{1:T}^{(b)}, y_{1:T,}$. Given the other parameters, beliefs about ρ can be updated from the autoregression:

$$\frac{1}{\sqrt{1-\gamma^{(b)}}} x_t^{(b)} = \rho \frac{1}{\sqrt{1-\gamma^{(b)}}} x_{t-1}^{(b)} + \sigma^{(b)} \eta_t.$$

   Define

$$\hat{\sigma}_\rho^2 \equiv [\sigma_\rho^{-2} + \frac{\Sigma_{s=2}^T (x_{s-1}^{(b)})^2}{(\sigma^{(b)})^2 (1-\gamma^{(b)})}]^{-1}$$
$$\hat{\mu}_\rho \equiv \hat{\sigma}_\rho^2 [\frac{\mu_\rho}{\sigma_\rho^2} + \frac{\Sigma_{s=2}^T x_{s-1}^{(b)} x_s^{(b)}}{(\sigma^{(b)})^2 (1-\gamma^{(b)})}]$$

   The posterior of ρ is $N(\hat{\mu}_\rho, \hat{\sigma}_\rho^2)$ and thus we draw $\rho^{(b+1)} \sim N(\hat{\mu}_\rho, \hat{\sigma}_\rho^2)$.

2. Draw $\mu^{(b+1)}|\rho^{(b+1)}, \sigma^{(b)}, x_{1:T}^{(b)}, y_{1:T,}$. Given the other parameters, beliefs about μ can be updated from the regression:

$$y_t - x_t = \mu + e_t.$$

Consider the matrix form

$$\begin{pmatrix} y_1 - x_1 \\ \vdots \\ y_T - x_T \end{pmatrix} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \mu + \begin{pmatrix} e_1 \\ \vdots \\ e_T \end{pmatrix}, \begin{pmatrix} e_1 \\ \vdots \\ e_T \end{pmatrix} \sim N(0, \sigma^2 \underbrace{\begin{bmatrix} \gamma & & 0 \\ & \ddots & \\ & & \gamma \\ 0 & & \gamma + \Gamma^2 \end{bmatrix}}_{\Lambda})$$

Define

$$\hat{\sigma}_\mu^2 \equiv [\sigma_\mu^{-2} + \frac{T-1}{\gamma^{(b)}(\sigma^{(b)})^2} + \frac{1}{(\gamma^{(b)} + \Gamma^2)(\sigma^{(b)})^2}]^{-1}$$
$$\hat{\mu}_\mu \equiv \hat{\sigma}_\mu^2 [\frac{\mu_\mu}{\sigma_\mu^2} + \frac{\Sigma_{s=1}^{T-1}(y_s - x_s^{(b)})}{\gamma^{(b)}(\sigma^{(b)})^2} + \frac{y_T - x_T^{(b)}}{(\gamma^{(b)} + \Gamma^2)(\sigma^{(b)})^2}]$$

The posterior of μ is $N(\hat{\mu}_\mu, \hat{\sigma}_\mu^2)$ and thus we draw $\mu^{(b+1)} \sim N(\hat{\mu}_\mu, \hat{\sigma}_\mu^2)$.

3. Draw $\sigma^{(b+1)} | \mu^{(b+1)}, \rho^{(b+1)}, x_{1:T}^{(b)}, y_{1:T,}$. Given the other parameters, beliefs about μ can be updated from the regressions:

$$y_t - x_t = \mu + e_t$$
$$\frac{1}{\sqrt{1-\gamma^{(b)}}} x_t^{(b)} = \rho^{(b+1)} \frac{1}{\sqrt{1-\gamma^{(b)}}} x_{t-1}^{(b)} + \sigma^{(b)} \eta_t$$

Consider the matrix form of the first equation

$$\begin{pmatrix} y_1 - x_1 \\ \vdots \\ y_T - x_T \end{pmatrix} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \mu + \begin{pmatrix} e_1 \\ \vdots \\ e_T \end{pmatrix}, \begin{pmatrix} e_1 \\ \vdots \\ e_T \end{pmatrix} \sim N(0, \sigma^2 \underbrace{\begin{bmatrix} \gamma & & 0 \\ & \ddots & \\ & & \gamma \\ 0 & & \gamma + \Gamma^2 \end{bmatrix}}_{\Lambda})$$

Define

$$\hat{A} \equiv A + \frac{2T-1}{2}$$
$$\hat{B} \equiv B + \frac{\Sigma_{s=1}^{T-1}(y_s - x_s^{(b)} - \mu^{(b+1)})^2}{2\gamma^{(b)}} + \frac{(y_T - x_T^{(b)} - \mu^{(b+1)})^2}{2(\gamma^{(b)} + \Gamma^2)} + \frac{\Sigma_{s=2}^{T}(x_s^{(b)} - \rho^{(b+1)} x_{s-1}^{(b)})^2}{2(1-\gamma^{(b)})}$$

The posterior of σ^2 is $IG(\hat{A}, \hat{B})$ and thus we draw $(\sigma^{(b+1)})^2 \sim IG(\hat{A}, \hat{B})$.

4. Draw $x_{1:T}^{(b+1)}|\mu^{(b+1)}, \rho^{(b+1)}, \sigma^{(b+1)}, \gamma, y_{1:T,}$. This can be done using the standard Kalman filter and simulation smoother.

C Forecasting with Diagnostic Filters

C.1 General Model Setup

Define linear Gaussian state-space model

$$y_t = Z\alpha_t + \epsilon_t, \epsilon_t \sim N(0, H)$$
$$\alpha_{t+1} = T\alpha_t + R\eta_t, \eta_t \sim N(0, Q)$$
$$\alpha_1 \sim N(a_1, P_1), t = 1, ..., n.$$

Let Y_t denote the set of past observations $y_1, ..., y_t$.

C.2 Notation

Expression	
$a_{t\|t}$	$E(\alpha_t\|Y_t)$
a_{t+1}	$E(\alpha_{t+1}\|Y_t)$
$P_{t\|t}$	$Var(\alpha_t\|Y_t)$
P_{t+1}	$Var(\alpha_{t+1}\|Y_t)$
v_t	$y_t - Z_t a_t$
F_t	$Var(v_t\|Y_{t-1})$

C.3 Representative Function and Distorted Posterior

Define the representativeness of a state α at t as the likelihood ratio:

$$R_t(\alpha) = \frac{f(\alpha|Y_t)}{f(\alpha|Y_{t-1} \cup \{\alpha_{t|t-1}\})}$$

State α is more representative at t if the signal y_t received in this period raises the probability of that state relative to the case where the news equals the ex ante forecast, $y_t = Za_t$, as described in the denominator.

The forecaster then over-weighs representative states by using the distorted posterior,

$$f^\theta(\alpha_t|Y_t) = f(\alpha_t|Y_t)R_t(\alpha_t)^\theta \frac{1}{J_t},$$

where J_t is a normalization factor and θ is a constant.

C.4 Distorted Posterior Distribution in the Steady State

The distorted posterior $f^\theta(\alpha_t|Y_t)$ is normal in the steady state characterized by a time-varying mean $a^\theta_{t|t}$ and constant covariance Σ, where

$$a^\theta_{t|t} = a_{t|t} + \theta(a_{t|t} - a_t)$$
$$\Sigma = \bar{P} - \bar{P}Z'\bar{F}^{-1}Z\bar{P}$$

and $\bar{P}$ and $\bar{F}$ are the steady state solutions of P_{t+1} and F_t.

C.5 Proof of the Distorted Posterior Distribution in the Steady State

For $y_t = Za_t$, we have $a_{t|t} = a_t$ as we receive no news this case, i.e., $v_t = 0$. So $f(\alpha_t|Y_{t-1} \cup \{\alpha_{t|t-1}\}) \sim N(a_t, \Sigma)$.

As we now know the distributions of $f(\alpha_t|Y_{t-1} \cup \{\alpha_{t|t-1}\})$ and $f(\alpha_t|Y_t)$, the log likelihood of $f^\theta(\alpha_t|Y_t)$ is the following

$$\begin{aligned} lnf^\theta(\alpha_t|Y_t) &\propto -\frac{1}{2}(\alpha_t - a_{t|t})'\Sigma^{-1}(\alpha_t - a_{t|t}) - \frac{1}{2}\theta(\alpha_t - a_{t|t})'\Sigma^{-1}(\alpha_t - a_{t|t}) + \frac{1}{2}\theta(\alpha_t - a_t)'\Sigma^{-1}(\alpha_t - a_t) \\ &\propto -\frac{1}{2}\alpha_t'\Sigma^{-1}\alpha_t + \alpha_t'\Sigma^{-1}a_{t|t} - \frac{1}{2}\theta\alpha_t'\Sigma^{-1}\alpha_t + \theta\alpha_t'\Sigma^{-1}a_{t|t} + \frac{1}{2}\theta\alpha_t'\Sigma^{-1}\alpha_t - \theta\alpha_t'\Sigma^{-1}a_t \\ &\propto -\frac{1}{2}\alpha_t'\Sigma^{-1}\alpha_t + \alpha_t'\Sigma^{-1}(a_{t|t} + \theta(a_{t|t} - a_t)) \end{aligned}$$

By completing the square, we can see that $f^\theta(\alpha_t|Y_t)$ is normal with mean $a_{t|t} + \theta(a_{t|t} - a_t)$ and covariance Σ.

C.6 Implementation of the Diagnostic Filter

To forecast with diagnostic filters, we include an additional step in the previous Gibbs sampling procedure. We additonally draw $\alpha_{diag,t|t} \sim N(a_{t|t} + \theta(a_{t|t} - a_t), \Sigma)$, which is later used in the forecasting step and θ controls the degree of diagnosticity.

D Estimating Prior Beliefs of Forecasters

D.1 Forecasting with the Posterior Parameters and States

The algorithm described in the above section yields B samples of the posterior of the states and parameters of our model at each point in time t. We remove the first J samples as burn-in and re-index these samples by b as follows $\{\mu^{(b)}, \rho^{(b)}, \sigma^{(b)}, x_{1:t}^{(b)}\}_{b=1}^{B}$. We then use the following algorithm to produce a real-time forecast distribution for the housing growth rate at time t:

1. For each $b = 1, ..., B$, simulate a path of shocks $\{\eta_{t+h}^{(b)}\}_{h=1}^{H}$ from the standard normal distribution.

2. Starting from $h = 1$, construct a simulated path of the states over H subsequent periods using the following equations:
$$x_{t+h}^{(b)} = \rho^{(b)} x_{t+h-1|t}^{(b)} + \sqrt{1-\gamma^{(b)}}\sigma^{(b)}\omega_{t+h}^{(b)}$$

3. Use the simulated states to construct the forecast distribution of the quarterly rate $\{y_{t+h|t}^{(b)}\}$ where
$$y_{t+h|t}^{(b)} = \mu_{t|t}^{(b)} + x_{t+h|t}^{(b)}$$

4. The quarterly forecast of y_{t+h} given time t information is computed as
$$F_t y_{t+h} = \frac{1}{B}\Sigma_{b=1}^{B} y_{t+h|t}^{(b)}$$

5. The quarterly forecasts are then aggregated to annual forecasts by considering both the simulated forecasts and realized current year's quarterly rate.

D.2 Searching over Prior Beliefs

Let $\delta = (\mu_\mu, \sigma_\mu, \mu_\rho, \sigma_\mu, \Gamma)'$ be the set of parameters that we wish to optimize. Define the moment function as $m(\delta) = \sum \frac{Empirical Moment - Model Implied Moment}{Empirical Moment}$. The parameters are then estimated via the simulated method of moments (SMM) with an identity weighting matrix,

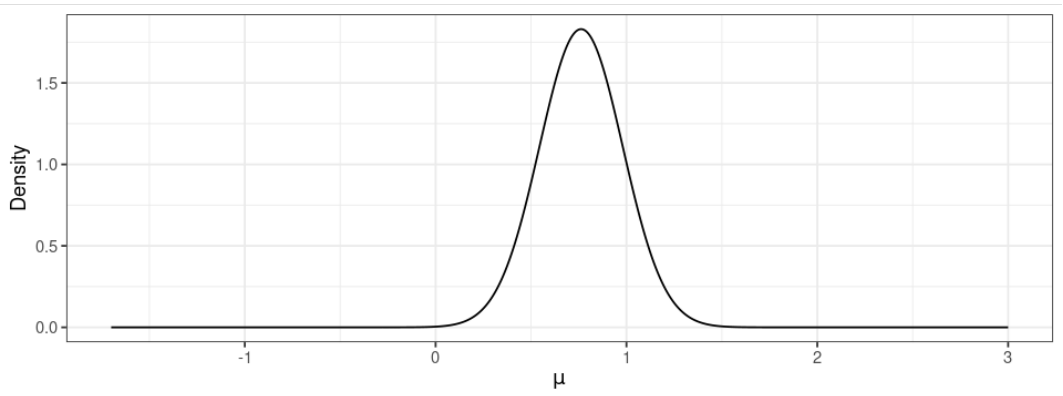
$$\hat{\Theta} = argmin_\delta m(\delta)' m(\delta).$$

Every evaluation of the moment function $m(\delta)$ requires us to add noise in the last observation and sample from the posterior of the model 30 times which corresponds to a simulated panel of 30 individuals. We use

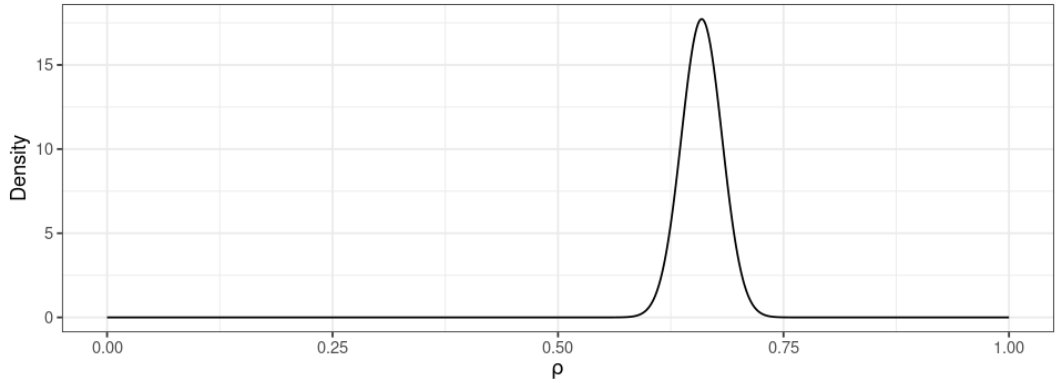
a burn-in sample of 2000 draws and keep the subsequent 8000 draws. The global minimum is found using the Bayesian optimization routine from the Python's scikit-optimize package.

E Additional Empirical Results

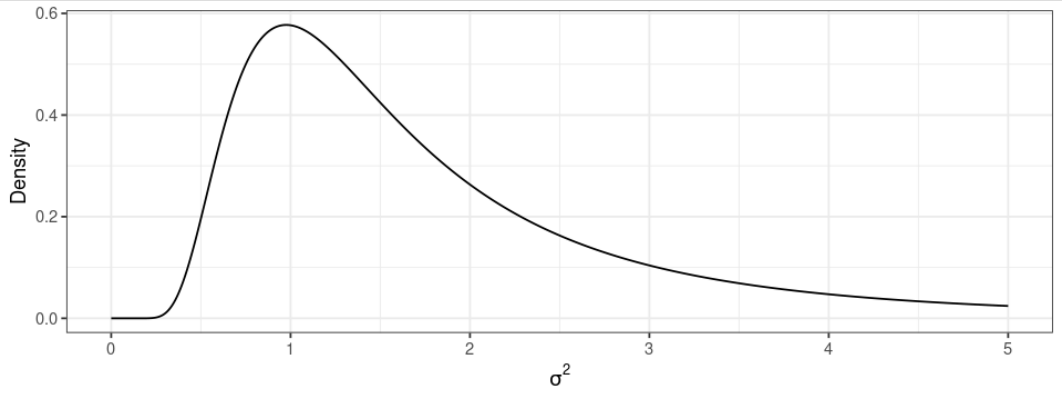
Figure A1: Prior Distributions of Parameters in the Benchmark Model



(a) Prior distribution of the long-run mean (μ)

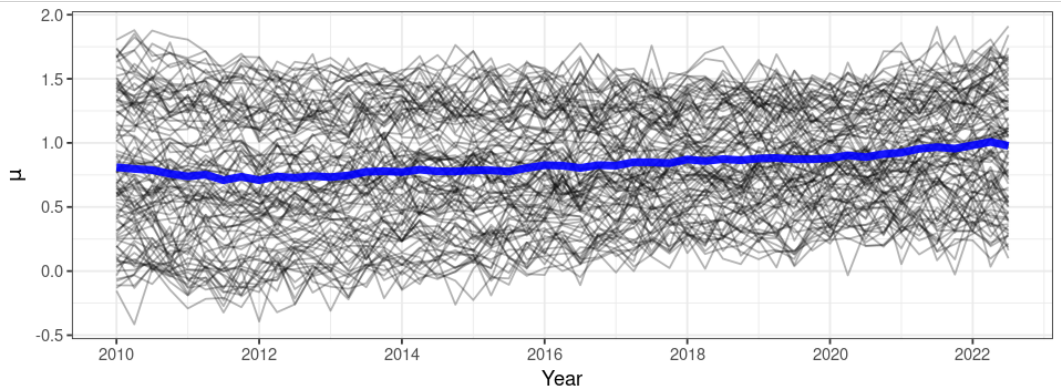


(b) Prior distribution of the persistent parameter (ρ)

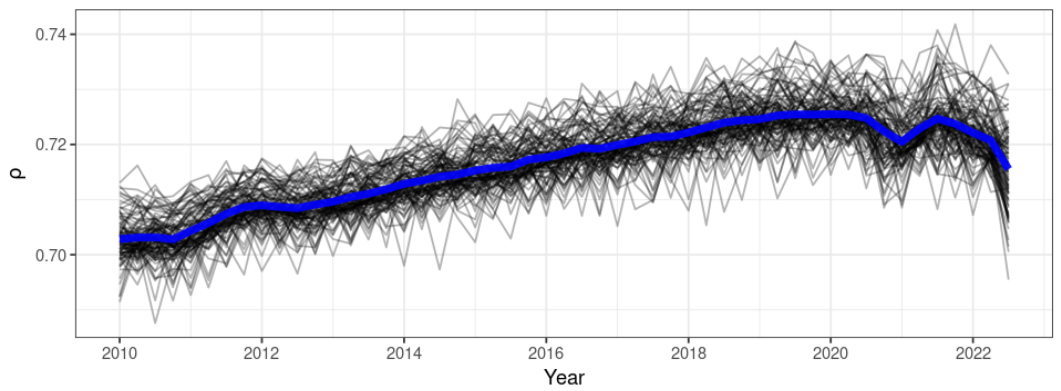


(c) Prior distribution of the volatility parameter (σ^2)

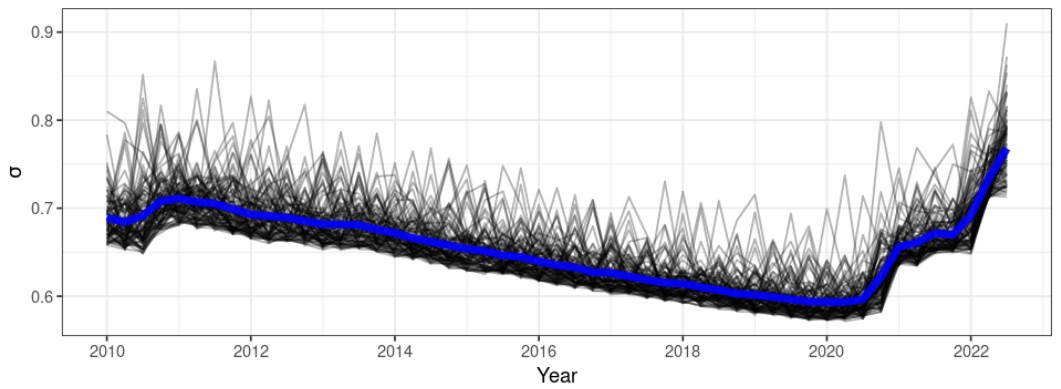
Figure A2: Posterior Distributions of Parameters in the Benchmark Model



(a) Posterior distribution of the long-run mean (μ)

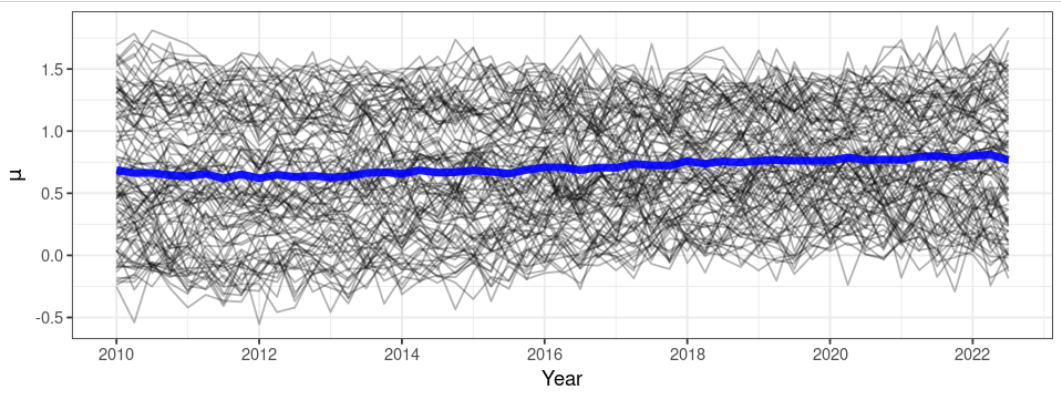


(b) Posterior distribution of the persistent parameter (ρ)

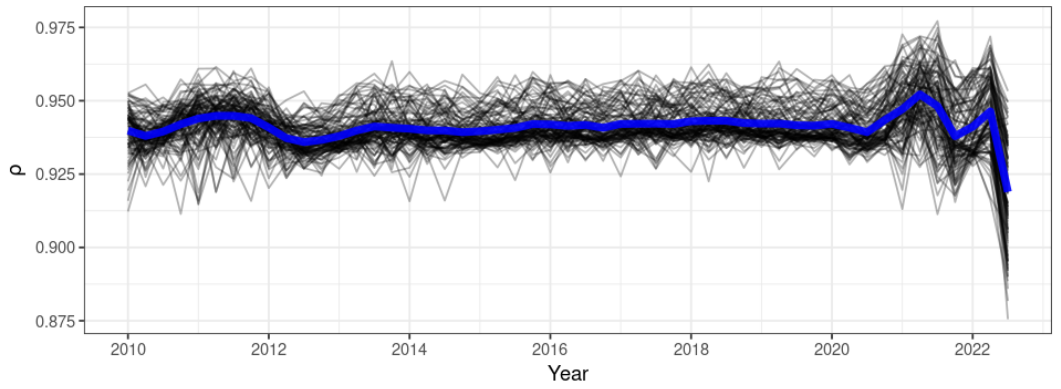


(c) Posterior distribution of the volatility parameter (σ^2)

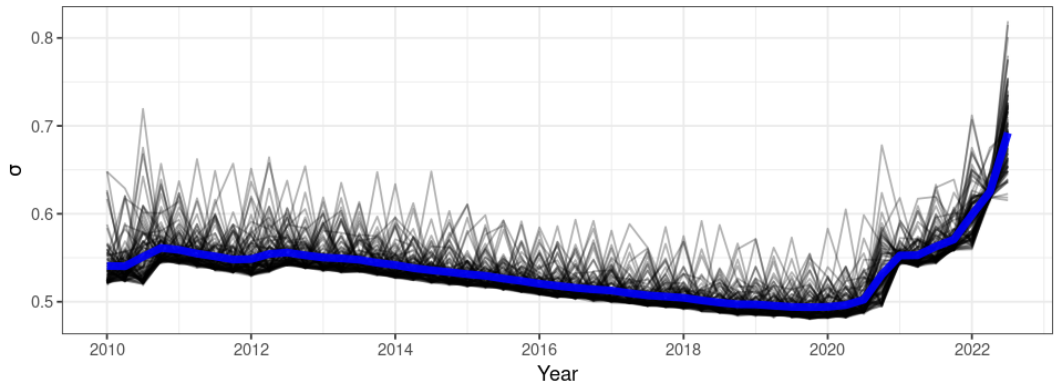
Figure A3: Posterior Distributions of Parameters in the No LR Model



(a) Posterior distribution of the long-run mean (μ)



(b) Posterior distribution of the persistent parameter (ρ)



(c) Posterior distribution of the volatility parameter (σ^2)

Table A1: Comparison of Residual Standard Errors from Individual Forecaster Regressions

Panel A: CG Regressions							
Moments	**Horizon**	**Data**	**Benchmark**	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
		(1)	(2)	(3)	(4)	(5)	(6)
CG	**1-Yr**	4.670	4.132	4.394	4.182	4.636	3.881
	2-Yr	4.057	3.666	3.738	3.678	4.290	3.341
	3-Yr	3.888	3.648	3.711	3.665	3.778	3.273
	4-Yr	3.942	3.810	3.866	3.822	3.607	3.418
Weighted Cost Function			**1.3**	**0.8**	**1.2**	**0.7**	**2.6**
Panel B: Sensitivity to Past House Price Growth							
Moments	**Horizon**	**Data**	**Benchmark**	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
Sensitivity to Past 1Q Growth	**1-Yr**	2.704	1.811	2.592	2.335	2.694	0.477
	2-Yr	2.240	1.786	1.946	1.839	2.214	0.112
	3-Yr	1.900	1.881	1.992	1.900	1.891	0.055
	4-Yr	1.756	1.933	2.039	1.954	1.767	0.048
Sensitivity to Past 2Y Growth	**1-Yr**	2.836	1.970	2.699	2.471	3.753	0.845
	2-Yr	2.262	1.792	1.951	1.845	2.757	0.189
	3-Yr	1.904	1.880	1.991	1.898	2.336	0.068
	4-Yr	1.757	1.932	2.038	1.953	2.147	0.049
Sensitivity to Past 4Y Growth	**1-Yr**	2.848	2.005	2.724	2.499	3.990	0.953
	2-Yr	2.265	1.786	1.946	1.839	2.935	0.209
	3-Yr	1.904	1.875	1.988	1.894	2.515	0.072
	4-Yr	1.756	1.929	2.035	1.950	2.285	0.049
Weighted Cost Function			**4.6**	**2.4**	**2.9**	**5.2**	**23.0**
Panel C: Biases and Inverse MZ Regressions							
Moments	**Horizon**	**Data**	**Benchmark**	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
Forecast Biases (%)	**1-Yr**	4.858	4.373	4.759	4.471	4.899	4.265
	2-Yr	4.136	3.697	3.784	3.716	4.548	3.361
	3-Yr	3.928	3.626	3.693	3.647	4.031	3.234
	4-Yr	4.031	3.826	3.884	3.835	3.757	3.419
Inverse MZ	**1-Yr**	2.595	2.130	2.813	2.577	4.431	1.156
	2-Yr	2.107	1.820	1.975	1.878	3.587	0.284
	3-Yr	1.835	1.883	1.994	1.903	3.014	0.090
	4-Yr	1.710	1.932	2.038	1.954	2.680	0.053
Weighted Cost Function			**2.5**	**1.7**	**1.8**	**6.3**	**9.3**
Panel D: Total Weighted Cost Function							
			Benchmark	**Overconfidence ($\phi = 2.25$)**	**Diagnostic Exp. ($\theta = 0.8$)**	**No LR.**	**No Disp.**
Total Weighted Cost Function			8.4	4.9	6.0	12.3	35.0

Notes: "Weighted Cost Function" for residual standard errors (RSE) of the regressions: |Model Implied RSE − Actual RSE|.

E.1 Parameter Identification

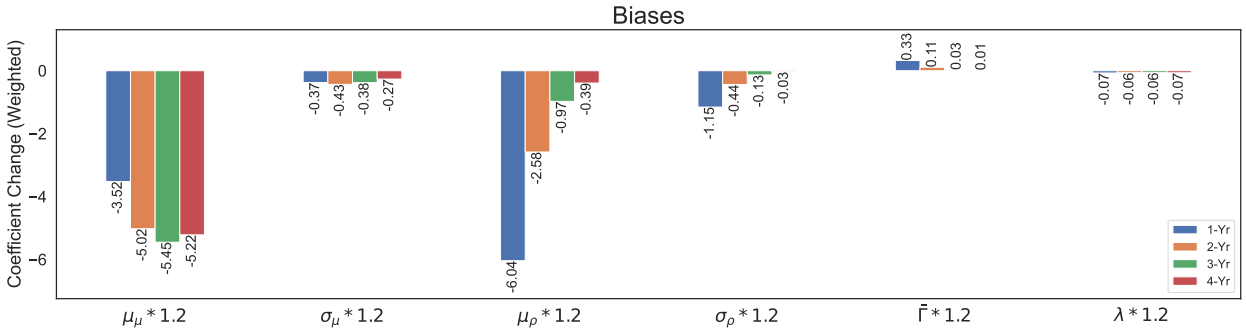
Inspired by Andrews, Gentzkow and Shapiro (2017), Figure A4 and Figure A5 vary each parameter by 20% up and down around its point estimate, and shows his this variation affects each of the moments used in estimation. If the moment changes substantially, the parameter in question is important for identification. Since the vertical axes are on the same scale across subplots, the variations are directly comparable. The graph is useful in providing intuition for the estimates.

The figure shows that increasing the prior mean of persistence, μ_ρ, affects several moments, but especially the sensitivity of house price forecasts to past price growth, and to a lesser extent the slope of the MZ regression, the slope of the CG regression, and the forecast dispersion. Since different moments respond in opposite directions, this parameter is well identified from the moments.

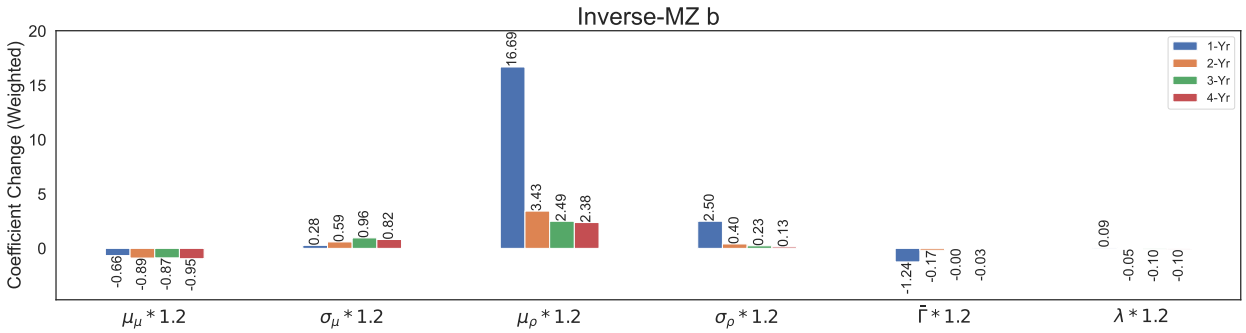
The prior mean of house price growth affects the mean forecast bias, the intercept of the MZ regression (not plotted), and the slope of the CG regression.

The figure also makes clear that the term structure of the forecasts is important in identifying the parameters. For example, the 3- and 4-year CG regression slopes are important for identifying the prior mean of ρ. The slope of the regression of forecasts on lagged realized house price growth, and how that slope varies across horizons, is also important for estimating the prior mean of ρ.

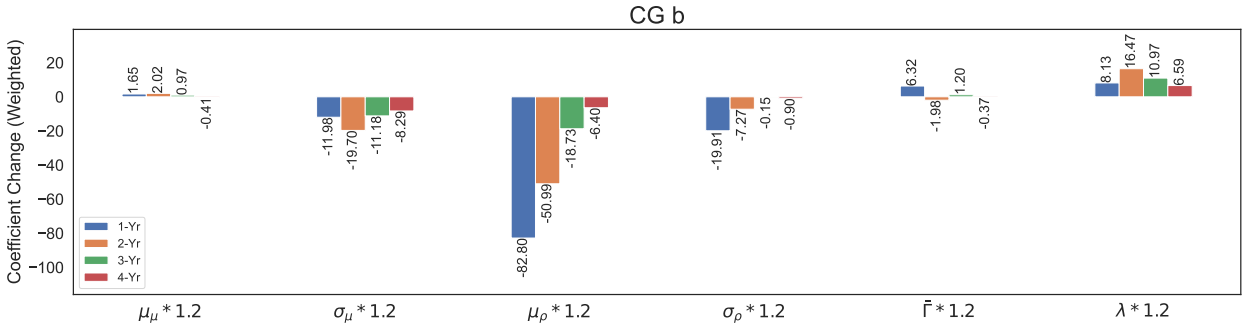
Figure A4: Sensitivity of Moments to Positive Perturbations of Parameters



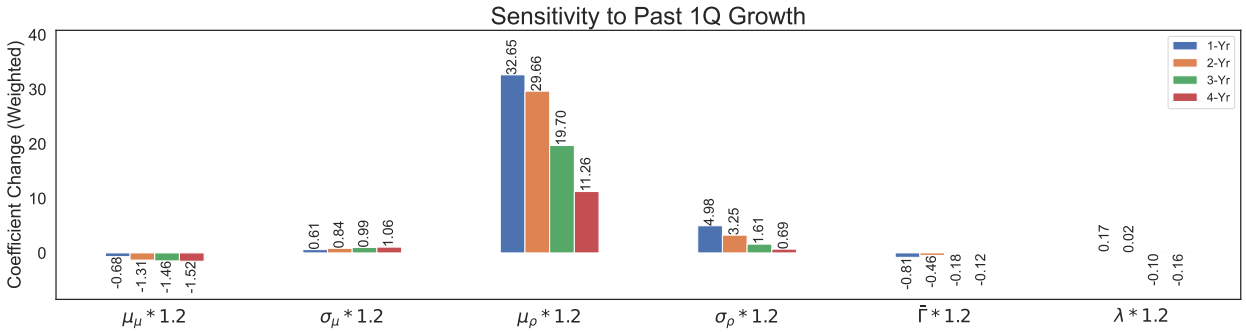
(a) Sensitivity of Forecast Biases to Parameter Perturbations



(b) Sensitivity of Slope of Inverse-MZ Regressions to Parameter Perturbations

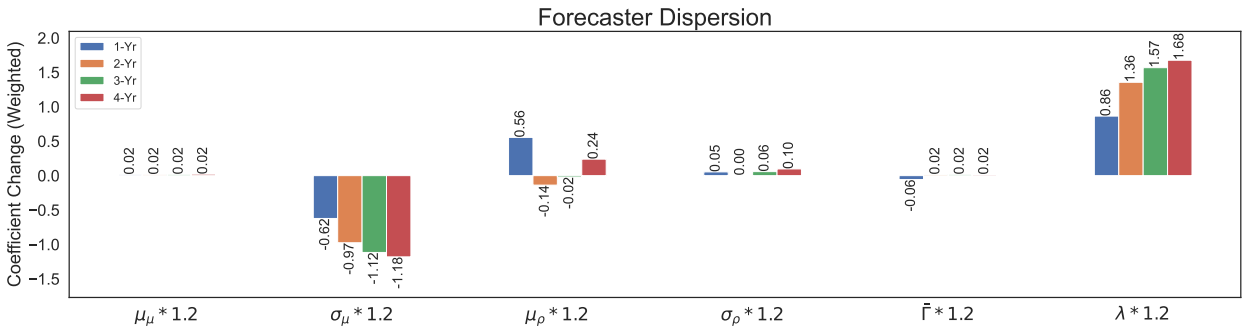


(c) Sensitivity of Slope of CG Regressions to Parameter Perturbations



(d) Sensitivity of Slope of Past 1Q Price Growth to Parameter Perturbations

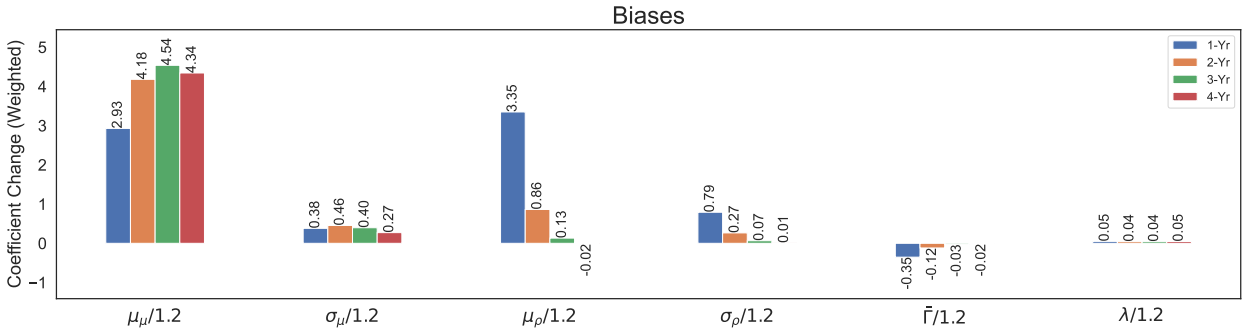
Figure A4: Sensitivity of Moments to Positive Perturbations of Parameters (cont.)



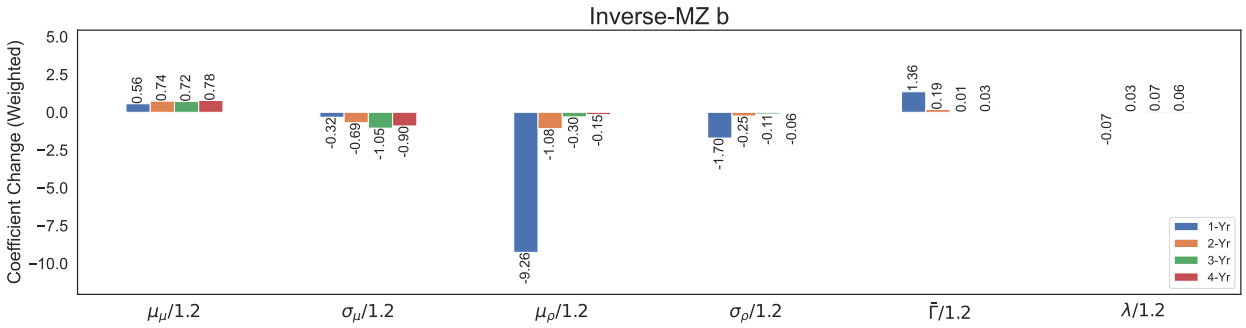
(e) Sensitivity of Dispersion to Parameter Perturbations

Notes: This figure shows how different empirical moments change if we increase the parameter value by 20% from its original point estimates. The changes in moment are calculated as the moment values based on perturbed parameters minus the moments based on the original parameter values, weighted by the weights in calculating the cost function. Each set of bars is for a particular moment. The different bars within each set (different colors) indicate the four different forecasting horizons.

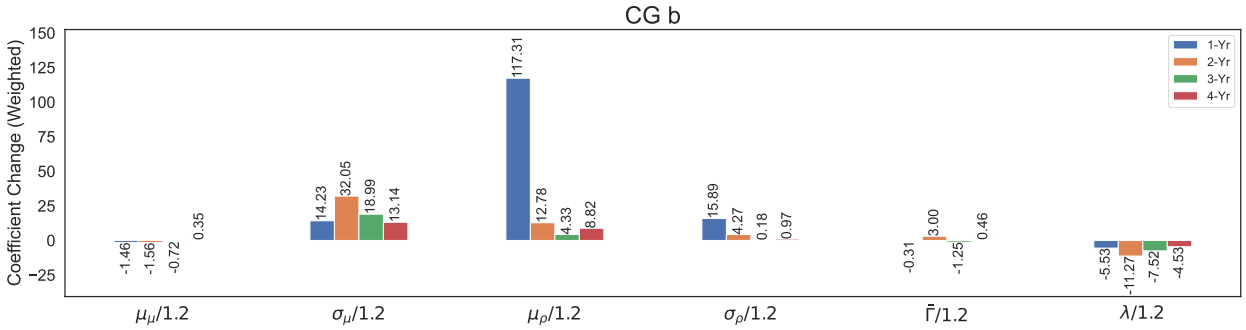
Figure A5: Sensitivity of Moments to Negative Perturbations of Parameters



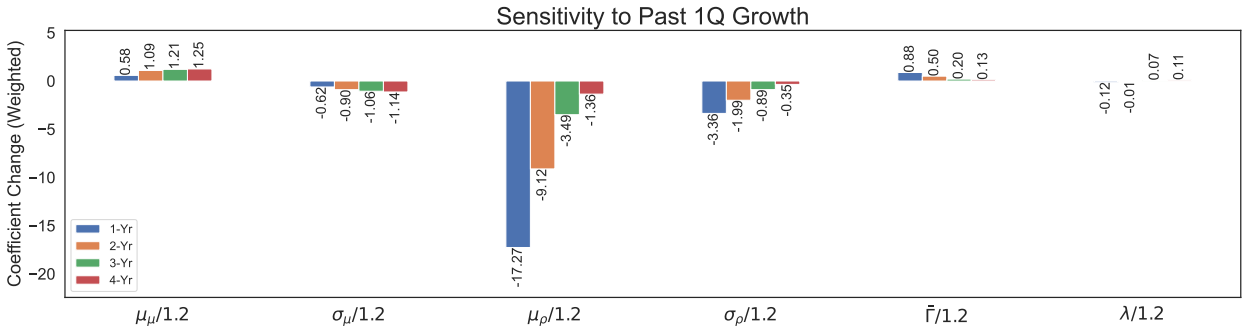
(a) Sensitivity of Forecast Biases to Parameter Perturbations



(b) Sensitivity of Slope of Inverse-MZ Regressions to Parameter Perturbations

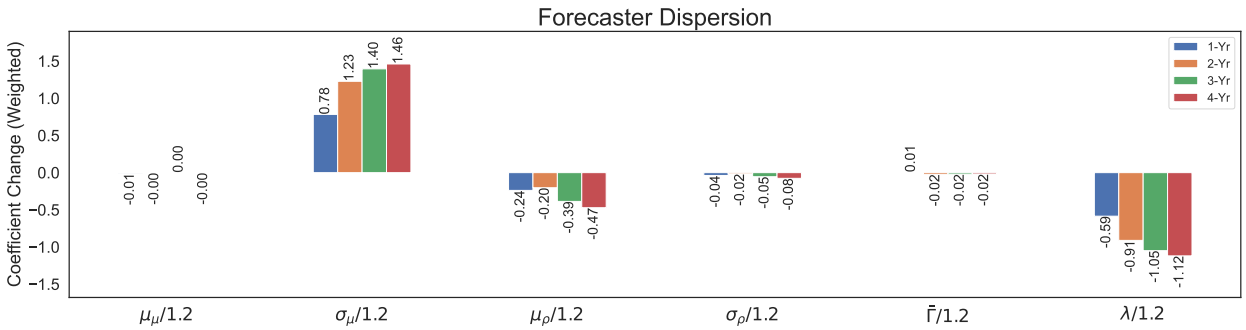


(c) Sensitivity of Slope of CG Regressions to Parameter Perturbations



(d) Sensitivity of Slope of Past 1Q Price Growth to Parameter Perturbations

Figure A5: Sensitivity of Moments to Negative Perturbations of Parameters (cont.)



(e) Sensitivity of Dispersion to Parameter Perturbations

Notes: This figure shows how different empirical moments change if we decrease the parameter value by 20% from its original point estimates. The changes in moment are calculated as the moment values based on perturbed parameters minus the moments based on the original parameter values, weighted by the weights in calculating the cost function. Each set of bars is for a particular moment. The different bars within each set (different colors) indicate the four different forecasting horizons.