

Asymmetric Overreaction in Household Inflation Expectations: The Role of News Tone

Mototsugu Shintani* Naoto Soma†

June 17, 2026

Abstract

This paper studies tone dependence in household inflation expectations. Using microdata from two large U.S. household surveys, the Michigan Survey of Consumers and the New York Fed Survey of Consumer Expectations, we find that household inflation expectations overreact more strongly to unfavorable news than to favorable news. The same asymmetry extends to the second moment of beliefs, with subjective inflation uncertainty rising after upward forecast revisions consistent with unfavorable news, and falling after downward revisions consistent with favorable news. To organize these facts, we develop an ambiguity-averse extension of diagnostic expectations in which households face ambiguity over the no-news reference distribution against which representativeness is evaluated.

1 Introduction

A growing literature on macroeconomic expectations finds that individual expectations overreact to news. Applying the forecast-error-revision regression of Coibion and Gorodnichenko (2015) at the individual level, Bordalo et al. (2020) first document overreaction in individual expectations among professional forecasters. They show that this pattern is well captured by the diagnostic-expectations (DE) framework of Bordalo, Gennaioli, and Shleifer (2018), in which news makes some states appear more representative, leading agents to overweight those states and generate overreaction. Since then, individual-level overreaction has been confirmed in additional surveys of professional forecasters and applied to a variety of macroeconomic and financial contexts (e.g., Bordalo, Gennaioli, La Porta, and Shleifer, 2019; Bianchi, Ilut, and Saijo, 2024a; L’Huillier, Singh, and Yoo, 2024). In all of these applications, however, the strength of overreaction is by construction the same across different types or directions of news. A natural question is whether overreaction depends on the *tone* of the news, that is, whether the news is favorable or unfavorable from the forecaster’s perspective. Whether overreaction is tone-dependent empirically, and how it would arise theoretically,

*Faculty of Economics, The University of Tokyo, Japan. Email: shintani@e.u-tokyo.ac.jp

†Faculty of Economics, Yokohama National University, Japan. Email: soma-naoto-wb@ynu.ac.jp

has received remarkably little attention.¹ This paper documents a news-tone asymmetry in U.S. household inflation expectations and develops a parsimonious theoretical framework in which the asymmetry arises endogenously from a single behavioral assumption.

We document two new empirical facts on household inflation expectations. The first fact concerns the first moment. Using the Michigan Survey of Consumers, where reported price-news tone is directly observed, we find that household inflation expectations overreact to both unfavorable and favorable price news, but the overreaction to unfavorable news is significantly stronger. The same asymmetry appears in the New York Fed Survey of Consumer Expectations (SCE). The SCE does not record reported news tone, so we use the sign of each household’s revision as a proxy, with upward revisions interpreted as unfavorable news. We find that upward revisions exhibit stronger overreaction than downward revisions.

The second fact concerns the second moment. Using the density forecasts in the SCE, we find that upward revisions in household point forecasts are accompanied by increases in subjective inflation uncertainty, while downward revisions are accompanied by decreases. The two facts therefore align in direction. Revisions in the unfavorable direction are accompanied by both stronger overreaction in point forecasts and higher subjective uncertainty, whereas revisions in the favorable direction are accompanied by the opposite pattern.

These two facts cannot both be reconciled within the canonical DE framework. In DE, agents form a Bayesian posterior in response to news and then compare it to a reference distribution describing what their beliefs would look like absent news. From this comparison, states whose probability has risen relative to the no-news reference distribution are viewed as representative of the news, and DE agents overweight them, producing overreaction. In the canonical formulation the no-news reference distribution is fixed and does not depend on the realized news, so the strength of overreaction is governed by a single parameter that is invariant to the direction of news, and the dispersion of the diagnostic posterior is invariant as well. Neither tone-dependent overreaction in the first moment nor the corresponding asymmetry in the second moment can therefore arise within this framework.

In this paper, we show that a simple modification of the canonical DE framework accounts for both facts at once. Under this modification, households face ambiguity over the appropriate no-news reference distribution, and conditional on the realized news tone, they select the reference distribution that is worst for themselves. We discipline the admissible set of reference distributions using a standard robust-control device (Hansen and Sargent, 2001, 2008; Hansen, Sargent, Turmuhambetova, and Williams, 2006). In particular, candidate reference distributions must lie within a relative-entropy ball around the reference distribution of canonical DE, containing those distributions that cannot be confidently distinguished from the canonical reference distribution on the basis of available data.

How does this single extension generate both empirical asymmetries? The mechanism operates through a single channel. Since the worst-case reference distribution differs across news tone, overreaction is amplified in one case and dampened in the other. In other words, the household’s pessimistic stance has different implications for different news. Under unfavorable news, the household compares against a tighter reference, treating the bad realization as a priori less likely and therefore especially informative. Against this tighter reference, the unfavorable news appears more representative, and overreaction is amplified. Under favor-

¹A notable exception is Aktuğ and Rezghi (2025), discussed below.

able news, it compares against a more diffuse reference, treating the good realization as a priori more likely and therefore less informative. Against the more diffuse reference, the favorable news appears less representative, and overreaction is dampened. The same tone-dependent choice of reference distribution explains the second-moment asymmetry. In our extension, as in canonical DE, subjective uncertainty equals the dispersion of the diagnostic posterior. This dispersion, however, is not fixed but tone-dependent because the dispersion of the chosen reference distribution itself depends on tone. After unfavorable news, the household compares against a tighter reference distribution, against which a broad range of states appears representative. The diagnostic posterior spreads out, and subjective uncertainty rises. After favorable news, the household compares against a more diffuse reference distribution, against which few states appear especially representative. The diagnostic posterior stays concentrated, and subjective uncertainty falls. The same tone-dependent choice of reference distribution therefore delivers both asymmetries through a single mechanism. The model nests canonical DE, which is recovered when households entertain no ambiguity over the reference distribution.

Our paper contributes to three strands of the literature. First, we contribute to the literature on expectation overreaction and diagnostic expectations. A large body of work studies overreaction in expectations using the forecast-error–revision approach of Coibion and Gorodnichenko (2015), beginning with professional forecasters in Bordalo et al. (2020), and subsequent empirical work documents overreaction using both survey data and laboratory experiments (Afrouzi et al., 2023; Duffy, Hanaki, and Yoo, 2025; Cornand, Hubert, and Portier, 2025). Bordalo, Gennaioli, and Shleifer (2018) introduce diagnostic expectations as a parsimonious model of representativeness-based overreaction, and subsequent work applies this framework to stock returns (Bordalo, Gennaioli, La Porta, and Shleifer, 2019), diagnostic bubbles (Bordalo, Gennaioli, Kwon, and Shleifer, 2021), macroeconomic forecasts (Bordalo et al., 2020), macroeconomic dynamics and business fluctuations (Bordalo, Gennaioli, and Shleifer, 2022; Bianchi, Ilut, and Saijo, 2024a), real credit cycles (Bordalo, Gennaioli, Shleifer, and Terry, 2026), asset prices and macro-finance (Maxted, 2024), and New Keynesian dynamics (L’Huillier, Singh, and Yoo, 2024). Closest within the diagnostic-expectations literature, Bianchi, Ilut, and Saijo (2024b) propose smooth diagnostic expectations, in which the strength of diagnostic distortions varies with the change in uncertainty between current and past beliefs. Across these settings, however, the no-news reference distribution against which representativeness is evaluated is either fixed or evolves independently of the realized tone of news, so diagnostic amplification is not itself chosen differently after favorable and unfavorable news. We extend this framework by letting the reference distribution depend on the realized news tone, while preserving the parsimony of the original formulation. The closest evidence on asymmetric overreaction is Aktuğ and Rezghi (2025), who study professional forecasts and find stronger overreaction to good news than to bad news. We instead study U.S. households, find stronger overreaction to unfavorable price news, and show that the same tone asymmetry extends to subjective uncertainty.

Second, we contribute to the literature on household inflation expectations, information exposure, news, and subjective uncertainty. Early work emphasizes the role of information diffusion, inflation news, and the media in household expectation formation (Carroll, 2003; Pfajfar and Santoro, 2013; Lamla and Lein, 2014; Dräger and Lamla, 2017; Larsen, Thorsrud, and Zhulanova, 2021). A related experimental and microdata literature shows

that household inflation expectations respond to statistical information, policy communication, and personal experience. Cavallo, Cruces, and Pérez-Truglia (2017) study learning from inflation statistics and supermarket prices. Coibion, Gorodnichenko, and Weber (2022) and D’Acunto, Hoang, and Weber (2022) study how monetary and fiscal communication shape household inflation expectations. D’Acunto, Malmendier, Ospina, and Weber (2021) show that households’ own grocery-price exposure affects inflation expectations, with positive price changes receiving more weight than negative price changes. Kuchler and Zafar (2019) show that personal experiences in housing and labor markets influence beliefs about aggregate outcomes. Weber et al. (2025) study learning in low- and high-inflation environments. D’Acunto et al. (2024) provide a recent overview of this literature. Our paper is especially close to recent work on inflation news. Chahrour, Shapiro, and Wilson (2025) study news selection and household inflation expectations, using the Michigan Survey news questions to show that exposure to higher-prices news raises inflation expectations whereas exposure to lower-prices news has little effect. Binder, Frank, and Ryngaert (2025) estimate the causal effect of television inflation news on household inflation expectations at a daily frequency and document stronger responses to positive than to negative CPI surprises. These papers establish that news exposure and news selection matter for household inflation expectations. We instead use reported price-news tone to measure tone-dependent overreaction through the individual forecast-error–revision regression, and we also show that the same asymmetry appears in the second moment of beliefs. By documenting that subjective uncertainty itself responds asymmetrically to news, we also contribute to research on subjective expectation uncertainty and household decisions, such as Ferman, Kuhnen, Li, and Ben-David (2024).

Third, our theory is related to the literature applying ambiguity aversion and robust control to macroeconomic belief formation. The robust-control approach of Hansen and Sargent (2001, 2008) and Hansen, Sargent, Turmuhambetova, and Williams (2006) provides the relative-entropy apparatus that we use to discipline deviations from the canonical reference distribution. Macroeconomic applications show how ambiguity aversion can generate pessimistic beliefs, confidence fluctuations, and endogenous amplification (Ilut and Schneider, 2014; Ilut and Saijo, 2021; Bhandari, Borovička, and Ho, 2025). Recent work also links ambiguity aversion directly to expectation bias, signal sensitivity, and household belief heterogeneity (Huo, Pedroni, and Pei, 2024; Michelacci and Paciello, 2024). Most closely, Baqaee (2020) microfounds asymmetric household inflation expectations with ambiguity aversion, with households overweighting inflationary news that reduces purchasing power and underweighting disinflationary news. Our contribution is to integrate this ambiguity-averse logic with diagnostic expectations. Rather than applying robustness to the agent’s perceived law of motion or to signal quality directly, we apply a relative-entropy restriction to the no-news reference distribution used for diagnostic updating. This yields a tone-dependent diagnostic posterior and jointly accounts for asymmetric overreaction in point forecasts and asymmetric subjective uncertainty.

The remainder of the paper is organized as follows. Section 2 documents the two empirical facts that motivate the analysis. Section 3 develops the ambiguity-averse diagnostic-expectations model. The final section concludes.

2 Empirical motivation

This section presents two empirical facts on household inflation expectations using the Michigan Survey of Consumers and the New York Fed Survey of Consumer Expectations. The first fact concerns the first moment of household beliefs, and the second concerns the second moment.

2.1 Data

2.1.1 Michigan Survey of Consumers

We use monthly microdata from the University of Michigan Surveys of Consumers, a nationally representative survey of U.S. households. The survey follows a rotating panel design: after a first interview, respondents are randomly selected to be re-interviewed six months later. In more recent years, a small subset of households are interviewed a third time six months after the re-interview. Our sample spans January 1980 to January 2024.

We focus on inflation expectations because inflation is the only macroeconomic variable in the Michigan Survey for which households report a numerical forecast. Expectations of other variables, such as unemployment, interest rates, and business conditions, are recorded only as directional responses (e.g., “go up,” “stay the same,” “go down”), which are not suitable for the forecast-error–revision regressions used in this paper.

Inflation expectations are collected in two steps. First, respondents answer “*During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?*” (*PX1Q1*). Second, they report a number to “*By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?*” (*PX1Q2*). We use the microdata variable *PX1* (“Expected Change in Prices During the Next Year”), which applies the sign implied by *PX1Q1* to the numeric response in *PX1Q2*.

To characterize respondents’ recent information, we use the question “*During the last few months, have you heard of any favorable or unfavorable changes in business conditions?*” Respondents first answer yes or no. Those answering yes are then asked a follow-up question, “*What did you hear? (Have you heard of any other favorable or unfavorable changes in business conditions?)*”, and describe in their own words any recent news they have heard. Each mention is subsequently classified by the survey interviewer according to topic, such as prices, employment, or government, and favorability, with the first and second mentions recorded as the microdata variables *NEWS1* and *NEWS2*. If the favorability of a reported change is not immediately clear, the interviewer is instructed to ask the respondent directly whether the mentioned change would be favorable or unfavorable, and to record the respondent’s own classification. Reported tone in our analysis therefore reflects the household’s own assessment of favorability rather than an ex-post coding by the researcher. The news question is administered *before* the inflation expectation question, which limits the scope for reverse priming from previously reported expectations to subsequently reported news. We pool *NEWS1* and *NEWS2* in classifying price-news tone: a household is recorded as reporting unfavorable price news if either *NEWS1* or *NEWS2* is coded as unfavorable price-increase news, and similarly for favorable price news.

We focus on respondents who report a unique price-news tone at the second wave, ex-

cluding the 18 households (0.3 percent of price-news reporters) who report both favorable and unfavorable news. Within our sample, 74.0 percent of households report unfavorable price news and 26.0 percent report favorable price news.

The mapping from price-news direction to reported tone is very strong: among households reporting price-increase news, 96.2 percent classify that news as unfavorable, while among households reporting price-decrease news, 86.8 percent classify that news as favorable. Given this strong mapping, we treat price-increase news as unfavorable and price-decrease news as favorable in the analysis below.

Table 1 reports summary statistics for the Michigan estimation sample, broken down by the reported tone of price news at the second interview. Unfavorable price news is considerably more common, with 3,268 of the 4,693 re-interview observations (70 percent) report unfavorable news, reflecting the well-documented tendency of households to notice price increases more readily than price decreases. Households reporting unfavorable news have higher average inflation forecasts ($\bar{f}_{i1} = 4.54$) than those reporting favorable news ($\bar{f}_{i1} = 2.44$), consistent with the expectation that unfavorable price signals push forecasts upward. The conditional mean forecast error differs in sign across the two groups. Favorable-news households exhibit a positive mean forecast error ($\overline{FE}_i = 0.45$), indicating that their forecasts were on average too low relative to the realization, while unfavorable-news households exhibit a negative mean forecast error ($\overline{FE}_i = -1.61$), indicating forecasts that were on average too high. The combination of overprediction after unfavorable news and underprediction after favorable news is suggestive of households overreacting to the tone of recent price news.

Table 1: Summary statistics: Michigan Survey

	All		Favorable		Unfavorable	
	Mean	SD	Mean	SD	Mean	SD
FR_i	-0.27	4.39	-0.86	3.85	-0.01	4.58
FE_i	-0.98	3.90	0.45	3.08	-1.61	4.06
f_{i1}	3.90	3.71	2.44	2.92	4.54	3.84
Observations	4,693		1,425		3,268	

Notes: Subsamples are defined by the reported tone of price news at the second interview: *Favorable* denotes households reporting favorable (price-decrease) news; *Unfavorable* denotes households reporting unfavorable (price-increase) news. $FR_i = f_{i1} - f_{i0}$ is the forecast revision, $FE_i = \pi_{m_{i1}+12} - f_{i1}$ is the forecast error, and f_{i1} is the second-interview inflation forecast. All values are in percentage points.

2.1.2 NY Fed Survey of Consumer Expectations

We use monthly microdata from the Federal Reserve Bank of New York’s Survey of Consumer Expectations (SCE). The SCE is a nationally representative, internet-based rotating panel of U.S. household heads. In each month, a fresh cross-section enters and a similar number exits. Respondents typically remain in the panel for up to twelve months. Our sample spans July 2013 to December 2024.

Point inflation forecasts are collected in two steps. First, respondents answer “*Over the next 12 months, do you think that there will be inflation or deflation?*” (*Note: deflation is*

the opposite of inflation)” (*Q8v2*). Second, they report a number to “What do you expect the rate of [inflation/deflation] to be over the next 12 months? Please give your best guess.” (*Q8v2part2*). We use the signed point forecast obtained from these two responses.

Density inflation forecasts are collected by asking respondents “In your view, what would you say is the percent chance that, over the next 12 months... ” the rate of inflation will be in each of ten bins ranging from 12% or higher deflation to 12% or higher inflation, with the probabilities summing to 100 percent (*Q9*). The mean and variance of the implied density are available as the microdata variables *Q9_mean* and *Q9_var*.

The SCE complements the Michigan Survey in two ways. First, it provides a longer monthly household panel, which allows us to absorb respondent fixed effects in the SCE FE–FR exercise. Second, it collects density forecasts for one-year-ahead inflation, which allows us to study household subjective uncertainty.

Table 2 reports summary statistics for the SCE estimation samples, broken down by the sign of the monthly forecast revision. The upper panel reports variables used in the forecast-error–revision analysis. Downward and upward revisers are roughly balanced in number (35,925 and 34,459, respectively), with approximately one-third of observations exhibiting zero revisions. As in the Michigan Survey, the conditional mean forecast error is asymmetric across revision directions. Downward revisers have a mean forecast error close to zero ($\overline{FE}_{it} = -0.19$), while upward revisers exhibit a strongly negative mean forecast error ($\overline{FE}_{it} = -2.81$), indicating systematic overprediction after upward revisions. The lower panel reports the subjective-uncertainty variables. The conditional mean changes in uncertainty differ in sign across revision directions. $\overline{\Delta IQR}_{it}$ is negative for downward revisers (-0.94) and positive for upward revisers (0.46), and the same pattern holds for $\overline{\Delta VAR}_{it}$ (-4.39 vs. 1.31). This sign pattern indicates that subjective uncertainty tends to decrease when households revise their inflation forecasts downward and increase when they revise upward, suggesting a systematic link between the direction of forecast revisions and changes in subjective uncertainty.

Table 2: Summary statistics: Survey of Consumer Expectations

<i>Forecast errors and revisions</i>						
	All		Downward		Upward	
	Mean	SD	Mean	SD	Mean	SD
FR_{it}	-0.09	3.47	-3.12	3.04	2.97	2.89
FE_{it}	-1.31	4.60	-0.19	3.86	-2.81	4.83
f_{it}	4.05	4.47	2.87	3.60	5.61	4.67
Observations	105,442		35,925		34,459	

<i>Subjective uncertainty</i>						
	All		Downward		Upward	
	Mean	SD	Mean	SD	Mean	SD
ΔIQR_{it}	-0.19	2.98	-0.94	3.45	0.46	3.48
ΔVAR_{it}	-1.18	20.75	-4.39	24.32	1.31	24.71
Observations	108,041		36,610		35,158	

Notes: Subsamples are defined by the sign of the monthly forecast revision FR_{it} : *Downward* denotes downward (favorable) revisions ($FR_{it} < 0$); *Upward* denotes upward (unfavorable) revisions ($FR_{it} > 0$). Observations with zero revisions are included in the *All* column but not in either subsample column. $FE_{it} = \pi_{t+12} - f_{it}$ is the forecast error, f_{it} is the point inflation forecast, ΔIQR_{it} and ΔVAR_{it} are the consecutive-month changes in the interquartile range and variance of the household's density forecast. The forecast-error sample and the subjective-uncertainty sample differ slightly due to different trimming and data requirements; the respective sample sizes are reported separately. All values are in percentage points.

2.1.3 Time-series patterns

Figure 1 plots the time-series evolution of reported price-news tone in the Michigan Survey alongside measures of the inflation environment. The red bars above the zero line show the share of re-interview households reporting unfavorable (price-increase) news, and the blue bars below the zero line show the share reporting favorable (price-decrease) news at the second interview. The black solid line is realized year-over-year CPI inflation at the interview date (right axis). The purple dash-dot line is the cross-sectional mean inflation forecast among Michigan price-news reporters, and the green dashed line is the corresponding mean forecast from the SCE (available from 2013).

The figure reveals a strong co-movement between reported tone and the inflation cycle. The share of unfavorable price news rises during episodes of elevated inflation—the early 1980s, the oil-price-driven increases in 2005 and 2008, and the post-pandemic inflation surge of 2021–2023—and falls during disinflationary periods. Conversely, the share of favorable price news expands when inflation moderates. This co-movement confirms that the reported tone of price news reflects the actual price environment rather than idiosyncratic household sentiment, and motivates our use of reported tone as a proxy for the direction of news that households process.

Both mean inflation forecasts broadly track realized CPI inflation. The Michigan mean

forecast (purple), computed from the relatively small sample of price-news reporters in the re-interview panel, is volatile but tracks the cyclical movements of CPI. The SCE mean forecast (green), based on a much larger monthly cross-section, is considerably smoother and closely follows the Michigan forecast in the overlap period, suggesting that the two surveys capture similar expectations dynamics despite their different designs.




Figure 1: Reported price-news tone, realized CPI inflation, and mean inflation forecasts

Notes: The red (blue) bars show the share of Michigan re-interview households reporting unfavorable (favorable) price news at the second interview (left axis). The black solid line is year-over-year CPI inflation at the interview date (right axis). The purple dash-dot line is the mean inflation forecast among Michigan price-news reporters, and the green dashed line is the mean inflation forecast from the SCE (available from July 2013). The right axis is centered at 2.5 percentage points.

Figure 2 plots the monthly cross-sectional mean of the two subjective-uncertainty measures available in the SCE—the interquartile range and the variance of each respondent’s density forecast—alongside the mean point inflation forecast. Both uncertainty measures exhibit pronounced time variation. They are low and stable during the pre-pandemic period of anchored inflation expectations, rise sharply during the post-pandemic inflation surge of 2021–2022, and subsequently decline as inflation moderates. The close co-movement between subjective uncertainty and the mean inflation forecast is consistent with the view that household uncertainty about future inflation responds to the same macroeconomic developments that drive point forecasts. The interquartile range and the variance display nearly identical cyclical patterns, confirming that both measures capture the same underlying variation in subjective uncertainty.




Figure 2: Subjective inflation uncertainty and mean inflation forecast in the SCE

Notes: Panel (a) plots the monthly cross-sectional mean of the interquartile range (IQR_{it} , blue solid line, left axis) of each respondent's density forecast for one-year-ahead inflation. Panel (b) plots the monthly cross-sectional mean of the variance (VAR_{it} , blue solid line, left axis). In both panels, the red dashed line shows the mean point inflation forecast (f_{it} , right axis). The sample consists of consecutive-month respondent pairs in the SCE from July 2013 to December 2024.

2.2 Tone dependence in household overreaction

We document tone-dependent overreaction in household inflation expectations using two complementary surveys. We begin with the Michigan Survey, which directly collects the reported tone of recent price news, and then turn to the SCE, whose longer per-household panel allows us to control for household fixed effects. For both surveys, we trim observations in the extreme 1st and 99th percentiles of the point forecast, the forecast revision, and the forecast error.

2.2.1 Michigan Survey

Let f_{i0} denote household i 's inflation forecast at the first interview, and let f_{i1} denote the corresponding forecast at the second interview. Let m_{i0} and m_{i1} denote the year-month (e.g., 2010m6) of the first and second interviews, respectively, and let π_τ denote the realized year-over-year inflation rate in month τ , so that $\pi_{m_{i1}+12}$ is realized inflation over the twelve months following the second interview.² The forecast error at the second interview and the forecast revision between the two interviews are

$$FE_i \equiv \pi_{m_{i1}+12} - f_{i1}, \quad FR_i \equiv f_{i1} - f_{i0}.$$

We label the tone of price news reported at the second interview as

$$\text{Tone}_i \in \{U, F\},$$

²Inflation is measured as the 12-month change in the seasonally adjusted headline CPI (FRED series CPIAUCSL), used throughout the paper.

where U stands for *unfavorable* price news (news suggesting price increases) and F for *favorable* price news (news suggesting price decreases). The corresponding tone indicator is

$$U_i \equiv \mathbf{1}\{\text{Tone}_i = U\}.$$

Every household in the sample reports exactly one of the two tones, so U_i equals one for households reporting unfavorable news and zero for those reporting favorable news.

We begin with a symmetric forecast-error–revision regression that does not distinguish between tones:

$$FE_i = \gamma_{m_{i1}, \text{Tone}_i} + \beta FR_i + \varepsilon_i, \quad (1)$$

where $\gamma_{m_{i1}, \text{Tone}_i}$ is a second-wave year-month-by-tone fixed effect. A negative β indicates overreaction on average: a larger revision predicts a forecast error of the opposite sign, so the revised forecast moves too far in the direction of the revision.

The symmetric specification restricts the FE–FR slope to be common across tones. To test whether the strength of overreaction depends on the reported tone, we augment the regression with a tone interaction:

$$FE_i = \gamma_{m_{i1}, \text{Tone}_i} + \beta_1 FR_i + \beta_2 (U_i \cdot FR_i) + \varepsilon_i. \quad (2)$$

For a household reporting favorable news, $U_i = 0$ and the slope of the forecast error on the revision is β_1 . For a household reporting unfavorable news, $U_i = 1$ and the interaction term adds β_2 , so the slope is $\beta_1 + \beta_2$. The coefficient β_2 therefore measures how the slope differs between the two tones. A nonzero β_2 indicates that the strength of overreaction depends on the reported tone, and a negative β_2 means that overreaction is stronger for unfavorable news.

The year-month-by-tone fixed effect $\gamma_{m_{i1}, \text{Tone}_i}$ in (1)–(2) plays two roles.³ First, it absorbs any difference in the average forecast error between the two tone groups interviewed in the same second-wave month. Second, it absorbs a shift in the forecast target that is specific to the Michigan Survey. In the Michigan Survey, each interview collects inflation forecasts over the subsequent twelve months, so f_{i0} is the household’s forecast of inflation between m_{i0} and $m_{i0} + 12$, and f_{i1} is its forecast between m_{i1} and $m_{i1} + 12$. The revision $FR_i = f_{i1} - f_{i0}$ therefore mixes a revision in expectations with a shift in the forecast target. In the re-interview sample, respondents are scheduled to be interviewed again six months after the first interview, so for observations with the scheduled re-interview gap the second-wave year-month m_{i1} pins down the first-wave year-month $m_{i0} = m_{i1} - 6$, and hence the forecast targets at both waves. If the rolling-horizon component is largely common across households sharing the same second-wave year-month and reported tone, this common component is absorbed by the second-wave year-month-by-tone fixed effect.

Both specifications are estimated with and without household-level demographic controls (age, household income, sex, number of children and adults in the household, region, marital status, years of education, and home ownership status). Appendix A reports a complementary specification with household fixed effects, estimated on the three-wave subsample.

³To avoid notational collision we use γ with calendar-month subscripts for the fixed-effect intercepts throughout the empirical sections. The regression slopes β, β_1, β_2 and $\gamma, \gamma_1, \gamma_2$ introduced later are distinguished by the form of their subscripts.

Table 3 reports the OLS estimates. Standard errors are clustered by second-wave month. Columns (1) and (2) present the symmetric specification (1). The pooled slope β is negative and highly significant in both columns, establishing that household inflation expectations overreact to price news on average.

The symmetric specification, however, restricts the FE–FR slope to be common across tones. Columns (3) and (4) relax this restriction. The favorable-tone slope β_1 and the unfavorable-tone slope $\beta_1 + \beta_2$ are both negative and statistically significant at the 1 percent level, but the unfavorable-tone slope is substantially more negative. The pooled slope in columns (1)–(2) lies between the two tone-specific slopes, reflecting the fact that the symmetric specification averages over the stronger overreaction to unfavorable news and the weaker overreaction to favorable news. The interaction coefficient β_2 is negative and statistically significant, confirming that the tone asymmetry is present and robust to the inclusion of demographic controls.

Table 3: Forecast-error–revision regressions in the Michigan Survey

	Symmetric		Tone-dependent	
	(1)	(2)	(3)	(4)
Pooled (β)	−0.423*** (0.022)	−0.424*** (0.022)		
Favorable (β_1)			−0.313*** (0.035)	−0.309*** (0.036)
Unfavorable ($\beta_1 + \beta_2$)			−0.449*** (0.025)	−0.451*** (0.025)
Difference (β_2)			−0.136*** (0.040)	−0.142*** (0.043)
Demographics	No	Yes	No	Yes
Second-wave year-month×tone FE	Yes	Yes	Yes	Yes
Observations	4,308	4,308	4,308	4,308

Notes: The dependent variable is the forecast error $FE_i = \pi_{m_{i1}+12} - f_{i1}$. The estimation sample is restricted to households who report exactly one price-news tone at the second interview. In columns (1) and (2), the regressor is the forecast revision $FR_i = f_{i1} - f_{i0}$ and the reported coefficient is the pooled slope β from (1). In columns (3) and (4), the regressors are FR_i and its interaction with the reported-tone indicator $U_i = \mathbf{1}\{\text{Tone}_i = U\}$; the reported coefficients are the favorable-tone slope β_1 , the implied unfavorable-tone slope $\beta_1 + \beta_2$, and the interaction β_2 from (2). All columns absorb the second-wave year-month-by-tone fixed effect $\gamma_{m_{i1}, \text{Tone}_i}$. Demographic controls consist of age, household income, sex, number of children and adults in the household, region, marital status, years of education, and home ownership status. Standard errors clustered by second-wave month are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

2.2.2 SCE

We next examine whether this pattern is also present in the SCE. The longer SCE panel allows us to absorb permanent forecaster heterogeneity through household fixed effects. We proxy the direction of news the household has received using the sign of the household’s monthly forecast revision. An upward revision is interpreted as the household having received

news in the unfavorable (inflationary) direction, while a downward revision is interpreted as news in the favorable direction.

Let f_{it} denote household i 's 12-month-ahead inflation point forecast reported in month t , so that π_{t+12} , defined as in Section 2.2, is realized inflation over the twelve months following month t . Define the forecast error and the revision between consecutive monthly interviews as

$$FE_{it} \equiv \pi_{t+12} - f_{it}, \quad FR_{it} \equiv f_{it} - f_{i,t-1}.$$

We formalize the sign-based proxy above through the indicator

$$U_{it} \equiv \mathbf{1}\{FR_{it} > 0\},$$

which plays the same role as the reported-tone indicator U_i of the Michigan Survey, equal to one for upward revisions in the unfavorable direction.⁴

The estimating equation parallels the Michigan framework. We begin with the symmetric specification:

$$FE_{it} = \alpha_i + \kappa_{\text{tenure}(it)} + \tau_t + \beta FR_{it} + \varepsilon_{it}, \quad (3)$$

where α_i is a household fixed effect, $\kappa_{\text{tenure}(it)}$ is a panel-tenure fixed effect that controls for systematic shifts associated with the household's position in the rotating SCE panel, and τ_t is a month fixed effect that absorbs aggregate inflation dynamics common to all households in month t . As in the Michigan specification, a negative β indicates overreaction on average.

To allow the strength of overreaction to differ by revision direction, we augment the regression with the tone interaction:

$$FE_{it} = \alpha_i + \kappa_{\text{tenure}(it)} + \tau_t + \beta_1 FR_{it} + \beta_2 (U_{it} \cdot FR_{it}) + \varepsilon_{it}. \quad (4)$$

The downward (favorable-direction) slope is β_1 , and the upward (unfavorable-direction) slope is $\beta_1 + \beta_2$. As in the Michigan regression, a negative slope indicates overreaction in the corresponding direction. The interaction coefficient β_2 measures the tone asymmetry: stronger overreaction to unfavorable news corresponds to $\beta_2 < 0$.

Table 4 reports the OLS estimates with household, panel-tenure, and month fixed effects. Standard errors are clustered by household to allow for arbitrary serial correlation in ε_{it} for the same individual. Column (1) presents the symmetric specification (3). The pooled slope β is negative and highly significant, confirming that overreaction to price news is also present in the SCE.

Column (2) relaxes the symmetry restriction. The downward slope β_1 and the upward slope $\beta_1 + \beta_2$ are both negative and statistically significant at the 1 percent level, but the upward slope is substantially more negative. As in the Michigan Survey, the pooled slope in column (1) lies between the two direction-specific slopes. The interaction coefficient β_2 is negative and statistically significant, replicating the tone asymmetry documented in the Michigan Survey while additionally absorbing permanent forecaster heterogeneity through the household fixed effect. Taken together, the Michigan and SCE evidence establishes a robust empirical target for the theory. We summarize the finding as follows.

⁴Applying the same sign-based proxy to the Michigan Survey, in place of its reported tone, yields the same tone asymmetry. See Appendix B.

Empirical Fact 1 (Tone-dependent overreaction). *Household inflation expectations overreact to inflation news in both directions, but the overreaction to unfavorable news is significantly stronger than the overreaction to favorable news.*

Table 4: Forecast-error–revision regressions in the SCE

	Symmetric (1)	Tone-dependent (2)
Pooled (β)	−0.441*** (0.005)	
Favorable (β_1)		−0.252*** (0.011)
Unfavorable ($\beta_1 + \beta_2$)		−0.642*** (0.011)
Difference (β_2)		−0.391*** (0.021)
Household FE	Yes	Yes
Panel-tenure FE	Yes	Yes
Month FE	Yes	Yes
Observations	103,196	103,196

Notes: The dependent variable is the 12-month-ahead forecast error $FE_{it} = \pi_{t+12} - f_{it}$ constructed from the SCE point forecast. In column (1), the regressor is the monthly revision $FR_{it} = f_{it} - f_{i,t-1}$ and the reported coefficient is the pooled slope β from (3). In column (2), the regressors are FR_{it} and its interaction with the sign-based tone proxy $U_{it} = \mathbf{1}\{FR_{it} > 0\}$; the reported coefficients are the downward-revision (favorable) slope β_1 , the implied upward-revision (unfavorable) slope $\beta_1 + \beta_2$, and the interaction β_2 from (4). All specifications absorb household, panel-tenure, and month fixed effects. Standard errors clustered by household are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

2.3 Point-forecast revisions and household subjective uncertainty

The previous subsection studies overreaction in point forecasts. We now turn to the second moment of beliefs using the density forecasts in the SCE. We examine whether revisions in households’ inflation point forecasts comove with revisions in their subjective inflation uncertainty.

The SCE collects a density forecast for year-ahead inflation, from which we construct two household-level uncertainty measures, the interquartile range IQR_{it} and the variance VAR_{it} . For each uncertainty measure $Y_{it} \in \{IQR_{it}, VAR_{it}\}$, the dependent variable is the consecutive-month change

$$\Delta Y_{it} = Y_{it} - Y_{i,t-1}.$$

We use the same monthly revision FR_{it} and sign-based tone proxy U_{it} as in Section 2.2.2, and ask whether the change in subjective uncertainty comoves with the revision.

We begin with a symmetric specification that does not distinguish between revision directions:

$$\Delta Y_{it} = \alpha_i + \kappa_{\text{tenure}(it)} + \tau_t + \gamma FR_{it} + u_{it}, \quad (5)$$

where α_i is a household fixed effect, $\kappa_{\text{tenure}(it)}$ is a panel-tenure fixed effect, and τ_t is a month fixed effect. A positive γ means that subjective uncertainty moves in the same direction as the signed revision: rising when the household revises upward and falling when it revises downward. Under the canonical diagnostic expectations of Bordalo et al. (2020), the diagnostic posterior variance is constant and does not respond to news, so the symmetric specification would predict $\gamma = 0$.

To allow the strength of the comovement to differ by revision direction, we augment the regression with the tone interaction, using the absolute revision $|FR_{it}|$:

$$\Delta Y_{it} = \alpha_i + \kappa_{\text{tenure}(it)} + \tau_t + \gamma_1 |FR_{it}| + \gamma_2 (U_{it} \cdot |FR_{it}|) + u_{it}. \quad (6)$$

These fixed effects absorb the corresponding household-specific, tenure-specific, and month-specific drifts in subjective uncertainty. The change in subjective uncertainty moves with the absolute revision at slope γ_1 for downward (favorable-direction) revisions and $\gamma_1 + \gamma_2$ for upward (unfavorable-direction) revisions. A negative γ_1 means that subjective uncertainty falls when the household revises its point forecast downward, and a positive $\gamma_1 + \gamma_2$ means that subjective uncertainty rises when the household revises upward. Therefore, comovement of subjective uncertainty with the direction of the revision—rising with unfavorable revisions and falling with favorable revisions—corresponds to $\gamma_1 < 0$ and $\gamma_1 + \gamma_2 > 0$. We trim observations in the extreme 1st and 99th percentiles of the point forecast, the forecast revision, and the current and lagged values of the interquartile range and the variance. Standard errors are clustered two ways by household and month.

Table 5 reports the estimates. Columns (1) and (2) present the symmetric specification (5). The pooled slope γ is positive and highly significant for both uncertainty measures, establishing that subjective uncertainty comoves with the signed revision on average. This finding already departs from the prediction of canonical diagnostic expectations, in which subjective uncertainty is invariant to news.

Columns (3) and (4) allow the comovement to differ by revision direction. In both columns, γ_1 is negative and $\gamma_1 + \gamma_2$ is positive, with both statistically significant: uncertainty falls with downward (favorable) revisions and rises with upward (unfavorable) revisions. The comovement documented by the symmetric specification is therefore not driven by one direction alone but holds in both directions separately. We summarize the finding as follows.

Empirical Fact 2 (Comovement of point-forecast revisions and subjective uncertainty). *Upward revisions in households' inflation point forecasts are associated with increases in subjective inflation uncertainty, whereas downward revisions are associated with decreases in subjective uncertainty.*

Together, Empirical Fact 1 and Empirical Fact 2 constitute the empirical target that the theoretical model developed in the next section is designed to rationalize.

Table 5: Point-forecast revisions and subjective uncertainty in the SCE

	Symmetric		Tone-dependent	
	(1) ΔIQR_{it}	(2) ΔVAR_{it}	(3) ΔIQR_{it}	(4) ΔVAR_{it}
Pooled (γ)	0.339*** (0.009)	1.727*** (0.064)		
Favorable (γ_1)			-0.365*** (0.010)	-1.902*** (0.074)
Unfavorable ($\gamma_1 + \gamma_2$)			0.313*** (0.013)	1.546*** (0.094)
Household FE	Yes	Yes	Yes	Yes
Panel-tenure FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	105,965	105,965	105,965	105,965

Notes: The dependent variable is the consecutive-month change in the interquartile range (ΔIQR_{it} , columns 1 and 3) or the variance (ΔVAR_{it} , columns 2 and 4) of the household’s year-ahead inflation density forecast. In columns (1) and (2), the regressor is the signed monthly revision $FR_{it} = f_{it} - f_{i,t-1}$ and the reported coefficient is the pooled slope γ from (5). In columns (3) and (4), the regressors are the absolute revision $|FR_{it}|$ and its interaction with the sign-based tone proxy $U_{it} = \mathbf{1}\{FR_{it} > 0\}$; the reported coefficients are the favorable (downward-revision) slope γ_1 and the implied unfavorable (upward-revision) slope $\gamma_1 + \gamma_2$ from (6). All specifications use SCE survey weights and absorb household, panel-tenure, and month fixed effects. Standard errors clustered two ways by household and month are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

3 An Ambiguity-Averse Diagnostic-Expectations Model

This section develops an expectation-formation model designed to generate the two empirical asymmetries documented in Section 2, tone-dependent overreaction in point forecasts (Fact 1) and tone-dependent responses of subjective uncertainty (Fact 2), within a single, parsimonious belief-formation mechanism. The model combines diagnostic expectations with an ambiguity-driven, tone-dependent choice of the reference distribution used to evaluate news.

In diagnostic expectations, the strength of overreaction is governed by representativeness, which in turn depends on a “no-news” reference distribution that appears in the denominator of the representativeness ratio. In the standard formulation, this reference distribution is fixed ex ante and is independent of the news received. Our key departure is to make this reference distribution *endogenous to the realized news tone*. Households are ambiguous about the appropriate reference distribution and, conditional on the tone of the realized signal, select the reference distribution that is worst for them, in the spirit of robust decision-making under model uncertainty. This single departure delivers both empirical asymmetries jointly.

The mechanism can be summarized in plain words. When the realized signal is unfavorable, the household selects, among the reference distributions it cannot statistically distinguish from a benchmark, the one under which the same signal is most surprising, that

is, the tightest reference dispersion. A tighter reference distribution scales up the representativeness ratio and therefore amplifies the diagnostic distortion of beliefs. This generates two simultaneous effects, namely a larger upward revision in the point forecast (the first-moment asymmetry of Fact 1) and a higher subjective uncertainty around that revised forecast (the second-moment asymmetry of Fact 2). When the realized signal is instead favorable, the worst-case reference distribution is the loosest one within the ambiguity set, so the diagnostic distortion is correspondingly muted, generating both a smaller point-forecast revision and a lower subjective uncertainty.

Section 3.1 briefly reviews diagnostic expectations. Section 3.2 introduces the general ambiguity-averse formulation. Section 3.3 specializes the framework to a linear-Gaussian state-space environment and derives a diagnostic Kalman filter with ambiguous reference dispersion. Section 3.4 establishes the two main implications. Proposition 3.7 delivers tone-dependent overreaction in FE–FR slopes (the theoretical counterpart of Fact 1), and Proposition 3.8 delivers tone-dependent subjective uncertainty (the theoretical counterpart of Fact 2).

3.1 Diagnostic expectations: a brief review

Let x_t denote a latent state capturing the persistent component of the forecast object. We normalize x_t so that higher values correspond to worse outcomes from the forecaster’s perspective. Under this normalization, upward innovations in x_t are “unfavorable” news, whereas downward innovations are “favorable” news. The state evolves according to an AR(1) process,

$$x_t = \rho x_{t-1} + u_t, \quad u_t \sim \mathcal{N}(0, \sigma_u^2), \quad \rho \in (-1, 1). \quad (7)$$

Household i observes a noisy signal

$$s_t^i = x_t + \varepsilon_t^i, \quad \varepsilon_t^i \sim \mathcal{N}(0, \sigma_\varepsilon^2), \quad (8)$$

with ε_t^i i.i.d. across i and t , and independent of u_t . Let S_t^i denote household i ’s information set at t .

Under Bayesian updating, household i forms the posterior density $f(x_t | S_t^i)$, which serves as the undistorted starting point for the diagnostic-expectations framework. Diagnostic expectations distort this Bayesian posterior by overweighting states that become relatively more likely after the news, where the increase in likelihood is measured against a “no-news” reference distribution. Following the diagnostic-expectations approach, we define this reference distribution as the posterior the household would form upon observing a no-news signal, i.e., a signal equal to the Bayesian prior mean. Let $x_{t|t-1}^i$ be the Bayesian prior mean at t and set

$$\bar{f}^{\text{ref}}(x | S_{t-1}^i) \equiv f(x | S_{t-1}^i \cup \{x_{t|t-1}^i\}).$$

Representativeness is the likelihood ratio

$$R_t^i(x) = \frac{f(x | S_t^i)}{\bar{f}^{\text{ref}}(x | S_{t-1}^i)}. \quad (9)$$

Given a parameter $\theta \geq 0$, the diagnostic density is

$$f^\theta(x | S_t^i) = \frac{f(x | S_t^i) (R_t^i(x))^\theta}{\int f(x | S_t^i) (R_t^i(x))^\theta dx}. \quad (10)$$

When $\theta > 0$, the posterior overweights the representative states, inducing overreaction relative to Bayes in the direction of the surprise.

3.2 Set of admissible reference distributions: general setup

We now formalize the ambiguity-averse departure outlined in the section introduction. Formally, we allow the reference density to vary within a family indexed by $\omega \in \Omega$:

$$f^{\text{ref}}(x | S_{t-1}^i; \omega).$$

Given this reference distribution, define representativeness and diagnostic beliefs as

$$R_{t,\omega}^i(x) = \frac{f(x | S_t^i)}{f^{\text{ref}}(x | S_{t-1}^i; \omega)}, \quad f^\theta(x | S_t^i; \omega) = \frac{f(x | S_t^i) (R_{t,\omega}^i(x))^\theta}{\int f(x | S_t^i) (R_{t,\omega}^i(x))^\theta dx}.$$

To discipline the ambiguity set in a way consistent with the idea that the data may not pin down a unique reference benchmark, we impose a robust-control style relative-entropy restriction, following Hansen and Sargent (2001, 2008). For two densities g and p , let

$$D_{\text{KL}}(g \| p) \equiv \int g(x) \log \frac{g(x)}{p(x)} dx$$

denote the Kullback–Leibler divergence (hereafter, KL divergence). For a radius $\delta \geq 0$, the set of admissible reference distributions consists of those that are statistically hard to distinguish from $\bar{f}^{\text{ref}}(\cdot | S_{t-1}^i)$, in the sense that

$$\Omega_t^i(\delta) \equiv \{\omega \in \Omega : D_{\text{KL}}(f^{\text{ref}}(\cdot | S_{t-1}^i; \omega) \| \bar{f}^{\text{ref}}(\cdot | S_{t-1}^i)) \leq \delta\}. \quad (11)$$

Let $\nu_t^i \in \{U, F\}$ denote the tone of the realized signal at date t for household i , where U represents unfavorable news (e.g., an upward inflation surprise) and F represents favorable news (e.g., a downward inflation surprise). Formally, ν_t^i is determined by the realized signal s_t^i ; in the Gaussian illustration below, it is given by the sign of the innovation. Conditional on the realized signal, the household selects the reference distribution that is worst for it within $\Omega_t^i(\delta)$. Since higher values of x_t are undesirable under our normalization, we model this worst-case selection as choosing the ω that maximizes the expected state under the resulting diagnostic density:

$$\omega_t^{i*}(S_t^i) \in \arg \max_{\omega \in \Omega_t^i(\delta)} \mathbb{E}_{f^\theta(\cdot | S_t^i; \omega)}[x_t]. \quad (12)$$

Beliefs used for forecasting are therefore

$$f^\theta(x | S_t^i; \omega_t^{i*}(S_t^i)). \quad (13)$$

The set $\Omega_t^i(\delta)$ does not depend on the news tone, but the worst-case choice $\omega_t^{i*}(S_t^i)$ does, through its dependence on the diagnostic density. In the Gaussian illustration of the next subsection, $\omega_t^{i*}(S_t^i)$ depends on S_t^i only through the tone ν_t^i , so we write the optimizer as $\omega_t^{i*}(\nu_t^i)$ in that special case.

Before turning to the Gaussian illustration, we note that both θ and δ must be positive for any tone asymmetry to arise.

Remark 3.1 (Ambiguity is irrelevant without diagnostic distortions). If $\theta = 0$, then $R_{t,\omega}^i(x)^\theta = 1$ for every $\omega \in \Omega_t^i(\delta)$, so $f^\theta(x | S_t^i; \omega) = f(x | S_t^i)$ for every admissible ω . Consequently, the maximization in (12) is immaterial and (13) equals the Bayesian posterior. Both $\theta > 0$ (diagnostic distortions) and $\delta > 0$ (ambiguity about the reference) are therefore required for the tone asymmetries derived below.

3.3 Gaussian illustration: amplified Kalman filter with ambiguous reference dispersion

To illustrate the implications of this framework, we derive the diagnostic updating rule in the linear-Gaussian state-space model and show how ambiguity over the dispersion of the reference distribution generates an asymmetric degree of overreaction to unfavorable and favorable news.

Let

$$x_{t|t-1}^i \equiv E[x_t | S_{t-1}^i], \quad P_{t|t-1}^i \equiv \text{Var}(x_t | S_{t-1}^i)$$

denote the Bayesian prior mean and variance. Define the Kalman gain and innovation as

$$K_t^i \equiv \frac{P_{t|t-1}^i}{P_{t|t-1}^i + \sigma_\varepsilon^2}, \quad \xi_t^i \equiv s_t^i - x_{t|t-1}^i. \quad (14)$$

Bayesian updating yields

$$x_{t|t}^i = x_{t|t-1}^i + K_t^i \xi_t^i, \quad P_{t|t}^i = (1 - K_t^i) P_{t|t-1}^i, \quad P_{t+1|t}^i = \rho^2 P_{t|t}^i + \sigma_u^2. \quad (15)$$

In the linear-Gaussian case, as shown by Bordalo et al. (2020), the standard diagnostic-expectations operator implies a scaled Kalman update. The innovation is amplified symmetrically by a factor $(1 + \theta)$,

$$x_{t|t}^{\theta,i} = x_{t|t-1}^i + (1 + \theta) K_t^i \xi_t^i. \quad (16)$$

We now bring the relative-entropy ambiguity set (11) into the linear-Gaussian illustration by characterizing the benchmark (no-news) reference distribution and then imposing the KL-ball restriction within a Gaussian reference distribution.

Lemma 3.2 (No-news benchmark reference in the linear-Gaussian model). *In the linear-Gaussian model, the no-news reference density $\bar{f}^{\text{ref}}$ defined in Section 3.1 satisfies*

$$\bar{f}^{\text{ref}}(x | S_{t-1}^i) = \mathcal{N}(x_{t|t-1}^i, P_{t|t}^i).$$

Proof. Under “no news,” the signal equals the Bayesian prior mean, $s_t^i = x_{t|t-1}^i$, so the innovation is $\xi_t^i = s_t^i - x_{t|t-1}^i = 0$. By the Kalman update (15), this implies $x_{t|t}^i = x_{t|t-1}^i$, while the posterior variance is $P_{t|t}^i$. Since the linear-Gaussian model yields a normal posterior, the resulting no-news posterior density is $\mathcal{N}(x_{t|t-1}^i, P_{t|t}^i)$. By definition of the no-news reference density, this is $\bar{f}^{\text{ref}}(x | S_{t-1}^i)$. $\square$

A naive reading of the no-news condition $\xi_t^i = 0$ might suggest that the reference variance should equal the prior variance $P_{t|t-1}^i$ rather than the posterior variance $P_{t|t}^i$. Lemma 3.2 clarifies why this intuition is incorrect. The reference distribution is the *posterior* obtained after Bayesian updating on the prior-mean signal, not the prior itself. Intuitively, even a signal that turns out to coincide with the prior mean is still informative in the sense that the household ex ante did not know whether it would coincide. Observing this realized null surprise sharpens the posterior relative to the prior. The no-news reference variance therefore equals the post-update variance $P_{t|t}^i$.

We next restrict the set of admissible reference distributions to Gaussian candidates that share the benchmark mean $x_{t|t-1}^i$ but have potentially different dispersion. The mean restriction is natural. Since the no-news benchmark mean is pinned down by the absence of innovation about x_t , ambiguity should apply not to the mean but to the dispersion of the reference distribution, which governs how surprising a given signal is relative to the no-news scenario. Write these candidates as

$$f_{t,\eta}^{\text{ref},i}(x) \equiv \mathcal{N}(x_{t|t-1}^i, \eta P_{t|t}^i), \quad \eta > 0.$$

In this Gaussian restriction, the general index ω from Section 3.2 reduces to the scalar dispersion-scale parameter η . The case $\eta = 1$ coincides with the benchmark (no-news) reference in Lemma 3.2. The case $\eta < 1$ implies a tighter reference-distribution dispersion and $\eta > 1$ implies a looser one. The relative-entropy bound then yields a simple interval for the admissible dispersion parameter.

Proposition 3.3 (Relative-entropy ambiguity over Gaussian references). *Fix household i and date t , and take the benchmark reference density to be $\bar{f}^{\text{ref}}(x | S_{t-1}^i) = \mathcal{N}(x_{t|t-1}^i, P_{t|t}^i)$. For $\eta > 0$, let $f_{t,\eta}^{\text{ref},i}(x) = \mathcal{N}(x_{t|t-1}^i, \eta P_{t|t}^i)$. Then*

$$D_{\text{KL}}\left(f_{t,\eta}^{\text{ref},i} \| \bar{f}^{\text{ref}}(\cdot | S_{t-1}^i)\right) = \frac{1}{2}(\eta - 1 - \log \eta).$$

Consequently, the KL-ball restriction $D_{\text{KL}}(f_{t,\eta}^{\text{ref},i} \| \bar{f}^{\text{ref}}(\cdot | S_{t-1}^i)) \leq \delta$ is equivalent to

$$\eta \in [\underline{\eta}(\delta), \bar{\eta}(\delta)],$$

where $\underline{\eta}(\delta) < 1 < \bar{\eta}(\delta)$ are the unique two solutions to $\frac{1}{2}(\eta - 1 - \log \eta) = \delta$ when $\delta > 0$ (and $\underline{\eta}(0) = \bar{\eta}(0) = 1$).

Proof. Fix i and t , and abbreviate $\mu \equiv x_{t|t-1}^i$ and $P \equiv P_{t|t}^i$. The benchmark density is

$$\bar{f}^{\text{ref}}(x | S_{t-1}^i) = \frac{1}{\sqrt{2\pi P}} \exp\left(-\frac{(x - \mu)^2}{2P}\right),$$

and the candidate density is

$$f_{t,\eta}^{\text{ref},i}(x) = \frac{1}{\sqrt{2\pi\eta P}} \exp\left(-\frac{(x-\mu)^2}{2\eta P}\right).$$

Therefore,

$$\log \frac{f_{t,\eta}^{\text{ref},i}(x)}{\bar{f}^{\text{ref}}(x \mid S_{t-1}^i)} = -\frac{1}{2} \log \eta - \frac{(x-\mu)^2}{2\eta P} + \frac{(x-\mu)^2}{2P} = -\frac{1}{2} \log \eta + \frac{(x-\mu)^2}{2P} \left(1 - \frac{1}{\eta}\right).$$

Taking expectations under $f_{t,\eta}^{\text{ref},i}$ and using $\mathbb{E}_{f_{t,\eta}^{\text{ref},i}}[(x-\mu)^2] = \eta P$ yields

$$D_{\text{KL}}\left(f_{t,\eta}^{\text{ref},i} \parallel \bar{f}^{\text{ref}}(\cdot \mid S_{t-1}^i)\right) = -\frac{1}{2} \log \eta + \frac{1}{2P} \left(1 - \frac{1}{\eta}\right) \eta P = \frac{1}{2}(\eta - 1 - \log \eta).$$

To establish the interval characterization, define $g(\eta) \equiv \eta - 1 - \log \eta$ for $\eta > 0$. Then

$$g'(\eta) = 1 - \frac{1}{\eta}, \quad g''(\eta) = \frac{1}{\eta^2} > 0,$$

so g is strictly convex on $(0, \infty)$ with unique minimizer at $\eta = 1$, where $g(1) = 0$. Moreover, $g(\eta) \rightarrow \infty$ as $\eta \downarrow 0$ and as $\eta \uparrow \infty$. Hence, for any $\delta > 0$, the equation $g(\eta) = 2\delta$ has exactly two solutions, one in $(0, 1)$ and one in $(1, \infty)$; denote them by $\underline{\eta}(\delta)$ and $\bar{\eta}(\delta)$. Strict convexity implies that the sublevel set $\{\eta > 0 : g(\eta) \leq 2\delta\}$ is the closed interval $[\underline{\eta}(\delta), \bar{\eta}(\delta)]$. When $\delta = 0$, the sublevel set reduces to $\{1\}$, so $\underline{\eta}(0) = \bar{\eta}(0) = 1$. $\square$

Motivated by Proposition 3.3, for notational simplicity we suppress the dependence of $\underline{\eta}$ and $\bar{\eta}$ on δ and interpret $[\underline{\eta}, \bar{\eta}]$ as $[\underline{\eta}(\delta), \bar{\eta}(\delta)]$ for some radius δ . Accordingly, we parameterize the family of admissible Gaussian reference densities by η as

$$f_{t,\eta}^{\text{ref},i}(x \mid S_{t-1}^i) \equiv \mathcal{N}(x_{t|t-1}^i, \eta P_{t|t-1}^i), \quad \eta \in [\underline{\eta}, \bar{\eta}]. \quad (17)$$

That is, the admissible reference distributions correspond to Gaussian “no-news” benchmarks that share the same mean $x_{t|t-1}^i$ but differ in dispersion through the scaling parameter η .

The next corollary records how this construction relates to the standard diagnostic-expectations specification of Bordalo et al. (2018, 2020). It corresponds to the special case $\delta = 0$ in our framework, providing a natural benchmark against which the tone asymmetries derived below can be compared.

Corollary 3.4 (Nesting standard diagnostic expectations). *If $\delta = 0$, then $\underline{\eta} = \bar{\eta} = 1$ and the set of admissible reference distributions collapses to a single benchmark density $\bar{f}^{\text{ref}}$. The ambiguity-averse diagnostic density (13) then coincides with the canonical diagnostic density. Hence $\delta > 0$ is necessary for any tone asymmetry to arise in the linear-Gaussian model.*

Proof. By Proposition 3.3, $\delta = 0$ implies $\underline{\eta}(0) = \bar{\eta}(0) = 1$, so $\eta \in [\underline{\eta}, \bar{\eta}]$ reduces to $\{1\}$. With $\eta = 1$, the candidate density (17) equals the benchmark $\bar{f}^{\text{ref}}(x \mid S_{t-1}^i)$ of Lemma 3.2, so the representativeness ratio (9) and the resulting diagnostic density (10) reduce to their standard form. $\square$

To prepare for the worst-case selection over the set of admissible reference distributions, we express the representativeness ratio and the diagnostic density as functions of the dispersion-scale parameter η . Let

$$R_{t,\eta}^i(x) = \frac{f(x | S_t^i)}{f_{t,\eta}^{\text{ref},i}(x | S_{t-1}^i)}, \quad f^\theta(x | S_t^i; \eta) = \frac{f(x | S_t^i) (R_{t,\eta}^i(x))^\theta}{\int f(x | S_t^i) (R_{t,\eta}^i(x))^\theta dx}.$$

The next proposition characterizes this diagnostic density in closed form. The key observation is that, in the linear-Gaussian environment, the diagnostic belief remains Gaussian, so it is fully summarized by its mean and variance. Both are explicit functions of the reference-dispersion parameter η , which allows us to study how worst-case selection of η shapes households' beliefs.

Proposition 3.5 (Diagnostic Kalman filter with ambiguous reference dispersion). *Assume $1 + \theta - \theta/\eta > 0$ (equivalently, $\eta > \theta/(1 + \theta)$). Then*

$$f^\theta(x | S_t^i; \eta) = \mathcal{N}(x_{t|t}^{\theta,i}(\eta), P_{t|t}^{\theta,i}(\eta)),$$

where

$$P_{t|t}^{\theta,i}(\eta) = \frac{P_{t|t}^i}{1 + \theta - \theta/\eta}, \quad (18)$$

and

$$x_{t|t}^{\theta,i}(\eta) = x_{t|t}^i + \lambda(\eta)(x_{t|t}^i - x_{t|t-1}^i), \quad \lambda(\eta) \equiv \frac{\theta}{\eta(1 + \theta) - \theta}. \quad (19)$$

Equivalently,

$$x_{t|t}^{\theta,i}(\eta) = x_{t|t-1}^i + (1 + \lambda(\eta))K_t^i \xi_t^i. \quad (20)$$

The standard diagnostic update (16) is nested. When $\eta = 1$, $\lambda(1) = \theta$.

Proof. By construction, the Bayesian posterior density is $f(x | S_t^i) = \mathcal{N}(x_{t|t}^i, P_{t|t}^i)$ and the reference density is $f_{t,\eta}^{\text{ref},i}(x | S_{t-1}^i) = \mathcal{N}(x_{t|t-1}^i, \eta P_{t|t}^i)$. Using (10) with the representativeness ratio built from $f_{t,\eta}^{\text{ref},i}$, the diagnostic density satisfies

$$f^\theta(x | S_t^i; \eta) \propto f(x | S_t^i) \left(\frac{f(x | S_t^i)}{f_{t,\eta}^{\text{ref},i}(x | S_{t-1}^i)} \right)^\theta = f(x | S_t^i)^{1+\theta} (f_{t,\eta}^{\text{ref},i}(x | S_{t-1}^i))^{-\theta}.$$

Let $\mu \equiv x_{t|t}^i$, $\mu_0 \equiv x_{t|t-1}^i$, and $P \equiv P_{t|t}^i$. Plugging in the Gaussian densities gives

$$\log f^\theta(x | S_t^i; \eta) = \text{const} - \frac{1 + \theta}{2P}(x - \mu)^2 + \frac{\theta}{2\eta P}(x - \mu_0)^2.$$

Expanding the squares and collecting terms yields

$$\log f^\theta(x | S_t^i; \eta) = \text{const} - \frac{1}{2P} \left[(1 + \theta - \theta/\eta)x^2 - 2((1 + \theta)\mu - (\theta/\eta)\mu_0)x + \text{const}' \right].$$

The leading coefficient is negative if and only if $1 + \theta - \theta/\eta > 0$, which is the integrability condition in the statement. Let $A \equiv 1 + \theta - \theta/\eta$ and $B \equiv (1 + \theta)\mu - (\theta/\eta)\mu_0$. Completing the square gives

$$f^\theta(x \mid S_t^i; \eta) \propto \exp\left(-\frac{A}{2P} \left(x - \frac{B}{A}\right)^2\right),$$

so $f^\theta(x \mid S_t^i; \eta) = \mathcal{N}(B/A, P/A)$. This implies

$$P_{t|t}^{\theta,i}(\eta) = \frac{P}{A} = \frac{P_{t|t}^i}{1 + \theta - \theta/\eta}.$$

Moreover,

$$x_{t|t}^{\theta,i}(\eta) = \frac{B}{A} = \frac{(1 + \theta)\mu - (\theta/\eta)\mu_0}{1 + \theta - \theta/\eta} = \frac{\eta(1 + \theta)\mu - \theta\mu_0}{\eta(1 + \theta) - \theta}.$$

Define $\lambda(\eta) \equiv \theta/(\eta(1 + \theta) - \theta)$. Then

$$\mu + \lambda(\eta)(\mu - \mu_0) = \frac{(\eta(1 + \theta) - \theta)\mu + \theta(\mu - \mu_0)}{\eta(1 + \theta) - \theta} = \frac{\eta(1 + \theta)\mu - \theta\mu_0}{\eta(1 + \theta) - \theta},$$

which shows $x_{t|t}^{\theta,i}(\eta) = x_{t|t}^i + \lambda(\eta)(x_{t|t}^i - x_{t|t-1}^i)$. Finally, substituting $x_{t|t}^i - x_{t|t-1}^i = K_t^i \xi_t^i$ from (15) yields (20). $\square$

3.4 Model implications for overreaction and subjective uncertainty

Equipped with the diagnostic Kalman filter under ambiguous reference dispersion, we now derive the implications that correspond to the two empirical exercises in Section 2. The first implication concerns the FE–FR regressions in Section 2.2. The model generates stronger overreaction after unfavorable news than after favorable news. The second implication concerns the SCE uncertainty regressions in Section 2.3. The model generates a rise in subjective uncertainty after unfavorable news and a fall after favorable news.

In the Gaussian environment, the realized news tone $\nu_t^i \in \{U, F\}$ introduced in Section 3.2 is determined by the sign of the Bayesian innovation:

$$\nu_t^i = U \iff \xi_t^i > 0 \quad \text{and} \quad \nu_t^i = F \iff \xi_t^i < 0.$$

The function $\lambda(\eta)$ is strictly decreasing for $\theta > 0$,

$$\lambda'(\eta) = -\frac{\theta(1 + \theta)}{(\eta(1 + \theta) - \theta)^2} < 0, \tag{21}$$

and therefore the effect of changing the reference dispersion depends on the sign of the innovation. Differentiating (19) gives

$$\frac{\partial x_{t|t}^{\theta,i}(\eta)}{\partial \eta} = \lambda'(\eta)(x_{t|t}^i - x_{t|t-1}^i) = \lambda'(\eta)K_t^i \xi_t^i. \tag{22}$$

If $\xi_t^i > 0$, increasing η lowers $x_{t|t}^{\theta,i}(\eta)$; if $\xi_t^i < 0$, increasing η raises $x_{t|t}^{\theta,i}(\eta)$, making the diagnostic revision less negative. Under ambiguity aversion, households choose η to maximize the expected, and thus worst-case, state. Intuitively, after unfavorable news ($\xi_t^i > 0$), the household selects the tightest admissible reference dispersion $\underline{\eta}$, which scales up the representativeness ratio and amplifies the diagnostic upward revision. After favorable news ($\xi_t^i < 0$), the household selects the loosest admissible reference dispersion $\bar{\eta}$, which dampens the diagnostic downward revision. The next proposition states this formally.

Proposition 3.6 (Worst-case reference dispersion is tone dependent). *Assume $\theta > 0$ and $1 + \theta - \theta/\underline{\eta} > 0$, and let the ambiguity set be $\eta \in [\underline{\eta}, \bar{\eta}]$. Under max-min reference selection,*

$$\eta_t^{i*}(\nu_t^i) \in \arg \max_{\eta \in [\underline{\eta}, \bar{\eta}]} x_{t|t}^{\theta,i}(\eta),$$

the worst-case choice satisfies

$$\eta_t^{i*}(U) = \underline{\eta}, \quad \eta_t^{i*}(F) = \bar{\eta}.$$

Consequently, the effective amplification terms satisfy

$$\psi_U \equiv \lambda(\underline{\eta}) > \lambda(\bar{\eta}) \equiv \psi_F \geq 0.$$

Proof. Since $\lambda'(\eta) < 0$, equation (22) implies that $x_{t|t}^{\theta,i}(\eta)$ is decreasing in η when $\xi_t^i > 0$ and increasing in η when $\xi_t^i < 0$. Therefore, when $\nu_t^i = U$ the maximizer is $\eta = \underline{\eta}$, and when $\nu_t^i = F$ the maximizer is $\eta = \bar{\eta}$. The inequality $\psi_U > \psi_F$ follows from $\underline{\eta} < \bar{\eta}$ and monotonicity of λ . Finally, $\psi_F \geq 0$ because $\lambda(\eta) = \theta/(\eta(1 + \theta) - \theta)$ is nonnegative on the admissible set under the integrability condition in Proposition 3.5. $\square$

Propositions 3.5–3.6 deliver tone-dependent diagnostic amplification. By (20), the diagnostic posterior mean under the worst-case reference distribution satisfies

$$x_{t|t}^{\text{AA},i}(\nu_t^i) \equiv x_{t|t}^{\theta,i}(\eta_t^{i*}(\nu_t^i)) = x_{t|t-1}^i + (1 + \psi_{\nu_t^i})K_t^i \xi_t^i, \quad \psi_U > \psi_F \geq 0,$$

where the superscript “AA” denotes the ambiguity-averse selection of the reference. The amplification factor $\psi_{\nu_t^i}$ governs how much the diagnostic forecast overshoots the Bayesian update, and therefore how much of the resulting forecast error is predictable from the revision itself. The next proposition translates this amplification into the FE–FR slope estimated in Section 2.2.

Proposition 3.7 (Tone-dependent overreaction in FE–FR regressions). *Let $u_{it}^\nu = K_t^i \xi_t^i$ denote the Bayesian update in tone state $\nu \in \{U, F\}$, so that by (20) the diagnostic forecast revision is*

$$FR_{it}^\nu = (1 + \psi_\nu)u_{it}^\nu.$$

Suppose the realization equals the Bayesian forecast plus an innovation e_{it}^ν orthogonal to the Bayesian update, $\text{Cov}(e_{it}^\nu, u_{it}^\nu) = 0$. Then the population coefficient in a tone-specific forecast-error–revision regression is

$$\beta_\nu = -\frac{\psi_\nu}{1 + \psi_\nu}.$$

Hence, $\psi_U > \psi_F \geq 0$ implies

$$\beta_U < \beta_F \leq 0.$$

In the standard-DE benchmark, $\delta = 0$ implies $\underline{\eta} = \bar{\eta} = 1$ and hence $\psi_U = \psi_F = \theta$. Therefore,

$$\beta_U = \beta_F = -\frac{\theta}{1 + \theta},$$

so standard DE generates symmetric overreaction but not tone-dependent FE–FR slopes.

Proof. Relative to the Bayesian forecast, the diagnostic forecast overshoots by $\psi_\nu u_{it}^\nu$. Therefore the forecast error is

$$FE_{it}^\nu = e_{it}^\nu - \psi_\nu u_{it}^\nu.$$

Using $FR_{it}^\nu = (1 + \psi_\nu)u_{it}^\nu$ and $\text{Cov}(e_{it}^\nu, u_{it}^\nu) = 0$,

$$\frac{\text{Cov}(FE_{it}^\nu, FR_{it}^\nu)}{\text{Var}(FR_{it}^\nu)} = \frac{\text{Cov}(e_{it}^\nu - \psi_\nu u_{it}^\nu, (1 + \psi_\nu)u_{it}^\nu)}{\text{Var}((1 + \psi_\nu)u_{it}^\nu)} = -\frac{\psi_\nu}{1 + \psi_\nu}.$$

The function $x \mapsto -x/(1 + x)$ is strictly decreasing for $x \geq 0$, so $\psi_U > \psi_F \geq 0$ implies $\beta_U < \beta_F \leq 0$. For the standard-DE statement, Corollary 3.4 gives $\underline{\eta} = \bar{\eta} = 1$ when $\delta = 0$. Substituting this into the amplification term yields $\psi_U = \psi_F = \theta$, and the equality of the two slopes follows immediately from the formula for β_ν . $\square$

Proposition 3.7 is the model’s direct counterpart to the tone-specific FE–FR regressions in Section 2.2. It shows that the tighter worst-case reference distribution after unfavorable news generates a more negative FE–FR slope in that direction, $\beta_U < \beta_F$, which corresponds to a negative interaction coefficient $\beta_2 < 0$ in regression (2). The standard-DE benchmark nests the symmetric case. It preserves overreaction when $\theta > 0$, but it cannot generate a difference between the two tone-specific slopes.

To derive the second-moment implication needed for the SCE uncertainty exercise, define the ambiguity-averse posterior variance by

$$P_{t|t}^{\text{AA},i}(\nu_t^i) \equiv P_{t|t}^{\theta,i}(\eta_t^{i*}(\nu_t^i)).$$

From (18),

$$\frac{\partial P_{t|t}^{\theta,i}(\eta)}{\partial \eta} = -\frac{P_{t|t}^i \theta}{\eta^2 (1 + \theta - \theta/\eta)^2} < 0 \quad (\theta > 0), \quad (23)$$

so the diagnostic posterior variance is strictly decreasing in the reference-dispersion parameter η . Combining (23) with Proposition 3.6 yields the next proposition.

Proposition 3.8 (Reference-relative subjective uncertainty). *Assume $\theta > 0$ and $1 + \theta - \theta/\underline{\eta} > 0$. Under the ambiguity set $\eta \in [\underline{\eta}, \bar{\eta}]$ and worst-case reference selection,*

$$P_{t|t}^{\text{AA},i}(U) = \frac{P_{t|t}^i}{1 + \theta - \theta/\underline{\eta}}, \quad P_{t|t}^{\text{AA},i}(F) = \frac{P_{t|t}^i}{1 + \theta - \theta/\bar{\eta}},$$

and therefore

$$P_{t|t}^{\text{AA},i}(U) > P_{t|t}^{\text{AA},i}(F).$$

If $\underline{\eta} < 1 < \bar{\eta}$, then

$$P_{t|t}^{\text{AA},i}(U) > P_{t|t}^i > P_{t|t}^{\text{AA},i}(F).$$

Equivalently, conditional on the Bayesian posterior variance $P_{t|t}^i$, the reference-relative uncertainty deviation

$$v_{it} \equiv P_{t|t}^{\text{AA},i}(\nu_t^i) - P_{t|t}^i$$

is a weakly increasing step function of the reference-relative diagnostic mean deviation

$$d_{it} \equiv x_{t|t}^{\text{AA},i}(\nu_t^i) - x_{t|t-1}^i.$$

In the standard-DE benchmark, $\delta = 0$ implies $\eta = 1$ in both tone states, so

$$P_{t|t}^{\text{AA},i}(U) = P_{t|t}^{\text{AA},i}(F) = P_{t|t}^i.$$

Thus the reference-relative uncertainty deviation v_{it} is identically zero under standard DE, and the subjective-uncertainty asymmetry in this proposition disappears.

Proof. By (23), $P_{t|t}^{\theta,i}(\eta)$ is strictly decreasing in η for $\theta > 0$. Proposition 3.6 implies $\eta_t^{i*}(U) = \underline{\eta}$ and $\eta_t^{i*}(F) = \bar{\eta}$, with $\underline{\eta} < \bar{\eta}$. Substituting these choices into (18) gives the stated formulas and the ordering $P_{t|t}^{\text{AA},i}(U) > P_{t|t}^{\text{AA},i}(F)$. Since $P_{t|t}^{\theta,i}(1) = P_{t|t}^i$, the comparison with the Bayesian benchmark follows from $\underline{\eta} < 1 < \bar{\eta}$. For the final statement, note that $d_{it} > 0$ if and only if $\nu_t^i = U$, while $d_{it} < 0$ if and only if $\nu_t^i = F$. Under $\underline{\eta} < 1 < \bar{\eta}$, the corresponding uncertainty deviations satisfy $v_{it} > 0$ in state U and $v_{it} < 0$ in state F . Thus, conditional on $P_{t|t}^i$, v_{it} is a weakly increasing step function of d_{it} . For the standard-DE benchmark, Corollary 3.4 implies that the ambiguity set collapses to $\eta = 1$. Substituting $\eta = 1$ into (18) gives $P_{t|t}^{\text{AA},i} = P_{t|t}^i$ in both tone states, so $v_{it} = 0$. $\square$

Proposition 3.8 compares the diagnostic posterior in tone states U and F relative to the no-news Bayesian benchmark. The SCE regression in Section 2.3, by contrast, uses survey-to-survey revisions of subjective uncertainty rather than deviations from a benchmark. The next proposition translates the result above into an implication for these net revisions, providing the form that can be confronted with the data.

Proposition 3.9 (Net mean revisions and subjective uncertainty revisions). *Let*

$$d_{it} \equiv x_{t|t}^{\text{AA},i}(\nu_t^i) - x_{t|t-1}^i$$

be the reference-relative diagnostic deviation of the posterior mean, and let

$$v_{it} \equiv P_{t|t}^{\text{AA},i}(\nu_t^i) - P_{t|t}^i$$

be the reference-relative deviation of subjective uncertainty. Suppose that, after conditioning on the Bayesian posterior variance scale, $v_{it} = g(d_{it})$ for a weakly increasing function g . Then

$$(d_{it} - d_{i,t-1})(v_{it} - v_{i,t-1}) \geq 0.$$

In particular,

$$d_{it} - d_{i,t-1} > 0 \quad \Rightarrow \quad v_{it} - v_{i,t-1} \geq 0,$$

and

$$d_{it} - d_{i,t-1} < 0 \quad \Rightarrow \quad v_{it} - v_{i,t-1} \leq 0.$$

In the standard-DE benchmark, $v_{it} = 0$ for all i, t . Hence $v_{it} - v_{i,t-1} = 0$ regardless of the sign of $d_{it} - d_{i,t-1}$, so the net-revision implication for subjective uncertainty is absent under standard DE.

Proof. Since g is weakly increasing, $d_{it} > d_{i,t-1}$ implies $g(d_{it}) \geq g(d_{i,t-1})$, and hence $v_{it} - v_{i,t-1} \geq 0$. The case $d_{it} < d_{i,t-1}$ is analogous and gives $v_{it} - v_{i,t-1} \leq 0$. Multiplying the two differences yields the stated inequality. The standard-DE statement follows from Corollary 3.4 and (18): when $\delta = 0$, the ambiguity set collapses to $\eta = 1$, so $v_{it} = 0$ in every period and the change in reference-relative uncertainty is identically zero. $\square$

We now bridge this theoretical result to the empirical specification of Section 2.3. We map the reported point forecast and density-implied uncertainty measure to model objects by $f_{it} = x_{t|t}^{\text{AA},i}(\nu_t^i)$ and $Y_{it} = P_{t|t}^{\text{AA},i}(\nu_t^i)$, identifying the household's reported beliefs with the diagnostic posterior under the worst-case reference dispersion. Under this identification, the reference-relative deviations d_{it} and v_{it} defined in Proposition 3.9 are not directly observed. The reported point forecast and uncertainty measure mix these reference-relative deviations with predictable trends in the underlying reference values themselves. Decompose the consecutive-month change in the reported uncertainty measure as

$$Y_{it} - Y_{i,t-1} = \underbrace{P_{t|t}^i - P_{t-1|t-1}^i}_{H_{it}^Y} + (v_{it} - v_{i,t-1}),$$

where H_{it}^Y is the reference-uncertainty drift. The empirical specification in Section 2.3 uses the absolute revision $|FR_{it}|$ interacted with the sign-based tone proxy U_{it} , together with household, panel-tenure, and month fixed effects. These fixed effects absorb the household-specific, tenure-specific, and month-specific drifts in H_{it}^Y . Similarly, decompose the survey-to-survey revision in the reported point forecast as

$$f_{it} - f_{i,t-1} = \underbrace{x_{t|t-1}^i - x_{t-1|t-2}^i}_{H_{it}^m} + (d_{it} - d_{i,t-1}),$$

where H_{it}^m is the reference-mean drift. Strictly speaking, the raw revision FR_{it} also contains H_{it}^m , so it does not isolate the reference-relative net revision $d_{it} - d_{i,t-1}$ that Proposition 3.9 directly concerns. Appendix C addresses this concern with a specification that removes H_{it}^m from the regressor by residualizing FR_{it} against the fixed effects, and confirms that the comovement is unchanged.

Proposition 3.9 predicts that subjective uncertainty falls with downward (favorable) revisions and rises with upward (unfavorable) revisions,

$$\gamma_1 < 0, \quad \gamma_1 + \gamma_2 > 0,$$

which is the pattern reported in Table 5. The standard-DE benchmark cannot generate this prediction. In that case, the model has no reference-relative uncertainty deviation in

the Gaussian specification, so the SCE net-revision implication for subjective uncertainty is absent.

We close the section with a comparative-statics result that ties the two empirical asymmetries to a single primitive, the ambiguity radius δ . Since δ governs how much room the household has to entertain alternative reference distributions, it also pins down the size of the tone-dependent reference distortion documented above.

Proposition 3.10 (Comparative statics in the ambiguity radius). *Assume $\theta > 0$ and $1 + \theta - \theta/\underline{\eta}(\delta) > 0$. The tone-dependent amplification gap and the reference-relative uncertainty gap are monotone in the ambiguity radius δ : as δ increases,*

$$\psi_U(\delta) - \psi_F(\delta) \quad \text{and} \quad P_{t|t}^{\text{AA},i}(U; \delta) - P_{t|t}^{\text{AA},i}(F; \delta)$$

are both strictly increasing. In particular, $\delta = 0$ implies no tone asymmetry (Corollary 3.4), while larger values of δ generate larger reference-relative asymmetries in both the first and the second moment of beliefs.

Proof. Strict convexity of $g(\eta) = \eta - 1 - \log \eta$ established in the proof of Proposition 3.3 implies that $\underline{\eta}(\delta)$ is strictly decreasing and $\bar{\eta}(\delta)$ is strictly increasing in δ , so the interval $[\underline{\eta}(\delta), \bar{\eta}(\delta)]$ widens strictly with δ . By Proposition 3.6, $\psi_U(\delta) = \lambda(\underline{\eta}(\delta))$ and $\psi_F(\delta) = \lambda(\bar{\eta}(\delta))$, and λ is strictly decreasing by (21); hence $\psi_U(\delta)$ is strictly increasing and $\psi_F(\delta)$ is strictly decreasing in δ , so their difference is strictly increasing. The argument for the variance ordering is analogous, using the strict monotonicity of $P_{t|t}^{\theta,i}(\eta)$ in η established by (23) in place of (21). At $\delta = 0$, Corollary 3.4 gives $\underline{\eta} = \bar{\eta} = 1$, so both gaps are exactly zero. $\square$

4 Conclusion

This paper documents a tone asymmetry in household inflation expectations along two dimensions of beliefs, and develops a parsimonious theory in which both arise from a single behavioral assumption. Using the Michigan Survey and the New York Fed SCE, we find that household inflation expectations overreact in both directions, with the overreaction significantly stronger following unfavorable than following favorable news. Density forecasts from the SCE reveal that the asymmetry extends to the second moment of beliefs, with subjective inflation uncertainty rising after upward revisions and falling after downward revisions. To organize these facts, we extend diagnostic expectations by endogenizing the no-news reference distribution against which representativeness is evaluated. Households face ambiguity over the appropriate reference distribution and, conditional on the realized news tone, select the reference distribution that is worst for them. The same tone-dependent choice of reference distribution delivers both asymmetries through a single mechanism. The radius of the ambiguity set parameterizes the strength of both asymmetries, and the model nests canonical diagnostic expectations at zero radius.

The tone asymmetry we document in household belief formation has the potential to deliver new implications for several macroeconomic questions, all of which involve translating an asymmetric belief response at the household level into asymmetries at the aggregate level. Since unfavorable price news triggers a more pronounced revision in household expectations,

aggregate inflation dynamics that follow symmetric shocks need not remain symmetric, with potential consequences for inflation persistence and the costs of disinflation. The simultaneous rise in subjective inflation uncertainty after unfavorable news, in turn, suggests that aggregate inflation uncertainty itself responds asymmetrically to news of different tone, which could feed back into household consumption and saving decisions. Both channels point to an asymmetric transmission of monetary policy. Disinflationary policy is particularly affected: monetary tightening conveys favorable inflation news, which moves household expectations less than unfavorable news would and may raise the cost of disinflation. Pursuing these implications quantitatively requires embedding the ambiguity-averse diagnostic-expectations mechanism developed here into a general-equilibrium environment, which would discipline the role of the ambiguity radius alongside other structural parameters. We leave this for future work.

A Three-wave households of the Michigan Survey with household and pair-tone fixed effects

This appendix presents a complementary specification using the subsample of households observed in three interview waves. Since three waves yield two current-wave revision observations per household, the specification can additionally include household fixed effects in the forecast-error-revision regression, while still controlling for the shift in the forecast target through a pair-tone fixed effect. The three-wave subsample is much smaller than the main two-wave sample used in Section 2.2, so the estimates here serve as a robustness check on the sign and ordering of the tone-specific slopes rather than as a primary specification.

For this appendix, we index the three interview waves by $s \in \{1, 2, 3\}$ and let m_{is} denote the calendar month of wave s . Let f_{is} denote household i 's inflation forecast reported at wave s , so that $\pi_{m_{is}+12}$, defined as in Section 2.2, is realized inflation over the twelve months following wave s . Since tone is now reported separately at each interview, we use wave-specific tone indicators.

$$\text{Tone}_{is} \in \{U, F\}, \quad U_{is} \equiv \mathbf{1}\{\text{Tone}_{is} = U\}.$$

The estimation sample restricts attention to wave- s observations for which $\text{Tone}_{is} \in \{U, F\}$. Since a revision requires a previous wave, the regression observations are indexed by $s \in \{2, 3\}$.

The forecast error at wave s and the revision between consecutive interviews are

$$FE_{is} = \pi_{m_{is}+12} - f_{is}, \quad FR_{is} = f_{is} - f_{i,s-1}.$$

Thus a single household contributes one observation at wave $s = 2$ (the revision between waves 1 and 2 evaluated under the wave-2 tone) and one observation at wave $s = 3$ (the revision between waves 2 and 3 evaluated under the wave-3 tone), provided the wave-2 and wave-3 tone restriction is satisfied.

The estimation equation parallels the main Michigan specification (2), with an added household fixed effect:

$$FE_{is} = \alpha_i + \psi_{m_{i,s-1}, m_{is}, \text{Tone}_{is}} + \beta_1 FR_{is} + \beta_2 (U_{is} \cdot FR_{is}) + \varepsilon_{is}, \quad s \in \{2, 3\}, \quad (24)$$

where α_i is a household fixed effect and $\psi_{m_{i,s-1}, m_{is}, \text{Tone}_{is}}$ is a fixed effect that interacts the previous-interview month, the current-interview month, and the current-wave tone. The pair-tone fixed effect plays the same role as in Section 2.2. It absorbs the shift in the forecast target that is common to households who share the same pair of consecutive interview months and the same current-wave tone. The household fixed effect α_i additionally absorbs permanent forecaster heterogeneity.

The estimation sample is small, with 282 three-wave households contributing 564 wave-2 and wave-3 observations, of which only 69 are favorable-tone observations. As a consequence, $\hat{\beta}_1$ is identified from limited variation and is estimated less precisely than $\hat{\beta}_1 + \hat{\beta}_2$. Table 6 reports the OLS estimates with standard errors clustered by household.

The implied unfavorable-tone slope $\beta_1 + \beta_2$ is negative and statistically significant at the 1 percent level, with a magnitude close to the OLS estimate in Table 3. The favorable-tone slope β_1 is also negative and significant at the 5 percent level, but estimated with a

Table 6: Forecast-error-revision regressions in the three-wave subsample of the Michigan Survey: household and pair-tone fixed effects

	(1)
	Three-wave households
Favorable (β_1)	−0.268** (0.117)
Unfavorable ($\beta_1 + \beta_2$)	−0.454*** (0.023)
Difference (β_2)	−0.186 (0.118)
Household FE	Yes
Pair×tone FE	Yes
Observations	564
Households	282
Favorable-tone observations	69

Notes: The estimation sample is restricted to households observed in three consecutive interview waves and who report exactly one price-news tone at each wave with a revision observation. The dependent variable is the forecast error $FE_{is} = \pi_{m_{is}+12} - f_{is}$. The regressors are the revision $FR_{is} = f_{is} - f_{i,s-1}$ and its interaction with the wave-specific tone indicator $U_{is} = \mathbf{1}\{\text{Tone}_{is} = U\}$. The reported coefficients are the favorable-tone slope β_1 , the implied unfavorable-tone slope $\beta_1 + \beta_2$, and the interaction β_2 . The specification absorbs household fixed effects and pair-tone fixed effects $\psi_{m_{i,s-1}, m_{is}, \text{Tone}_{is}}$ that interact the previous-interview month, the current-interview month, and the current-wave tone. Standard errors clustered by household are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

substantially wider standard error owing to the small number of favorable-tone observations. The interaction coefficient β_2 is therefore not statistically significant at conventional levels, although its sign matches the asymmetry documented in the main text. The three-wave evidence therefore supports the tone asymmetry documented in Section 2.2 in terms of the sign and ordering of the slopes, with the caveat that statistical precision on the asymmetry test is constrained by the limited number of favorable-tone observations.

B A sign-based news proxy in the Michigan Survey

The main specification in Section 2.2 uses the reported price-news tone, which the Michigan Survey records directly. In the model, the forecast revision $f_{i1} - f_{i0}$ equals the change in the household’s diagnostic posterior mean and therefore inherits the sign of the news innovation, so the sign of the revision is a model-consistent proxy for the news direction. This appendix verifies that the tone asymmetry is also obtained when the news direction is proxied by the sign of the revision alone, ignoring the reported tone, in the spirit of Aktuğ and Rezghi (2025) and paralleling the SCE specification of Section 2.2.2.

We re-estimate the Michigan regression while ignoring the reported tone, replacing the tone indicator U_i with the sign-based proxy $U_i^{\text{sign}} = \mathbf{1}\{FR_i > 0\}$ and absorbing a second-wave year-month fixed effect:

$$FE_i = \gamma_{m_{i1}} + \beta_1 FR_i + \beta_2 (U_i^{\text{sign}} \cdot FR_i) + \varepsilon_i.$$

Table 7 reports the estimates.⁵ The downward (favorable-direction) slope β_1 and the upward (unfavorable-direction) slope $\beta_1 + \beta_2$ are both negative and statistically significant, and the upward slope is substantially more negative than the downward slope. The tone asymmetry documented in Section 2.2 is therefore also present when the news direction is identified from the sign of the revision alone.

Table 7: Michigan forecast-error-revision regression with a sign-based news proxy

Downward (β_1)	−0.114*** (0.023)
Upward ($\beta_1 + \beta_2$)	−0.799*** (0.023)
Demographics	Yes
Second-wave year-month FE	Yes
Observations	4,489

Notes: The dependent variable is the forecast error $FE_i = \pi_{m_{i1}+12} - f_{i1}$. The reported tone is ignored: the regressors are the revision $FR_i = f_{i1} - f_{i0}$ and its interaction with the sign-based proxy $U_i^{\text{sign}} = \mathbf{1}\{FR_i > 0\}$, so the upward and downward slopes are identified from the sign of the revision alone, in the spirit of Aktuğ and Rezghi (2025). The reported rows are the downward-revision slope β_1 and the upward-revision slope $\beta_1 + \beta_2$. The specification absorbs a second-wave year-month fixed effect and includes demographic controls. Standard errors clustered by second-wave month are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

⁵The corresponding estimates without demographic controls are very similar and are omitted for brevity.

C Residualized revisions in the SCE uncertainty regression

As established in Section 3.4, Proposition 3.9 predicts a positive comovement between the reference-relative net revision $d_{it} - d_{i,t-1}$ and the reference-relative change in subjective uncertainty $v_{it} - v_{i,t-1}$. The empirical specification in Section 2.3 regresses the change in subjective uncertainty on the absolute revision $|FR_{it}|$ interacted with the sign-based tone proxy U_{it} , together with household, panel-tenure, and month fixed effects. As noted there, FR_{it} also contains the reference-mean drift H_{it}^m , so it does not isolate the theoretical object $d_{it} - d_{i,t-1}$ that Proposition 3.9 directly concerns. This appendix removes H_{it}^m from the regressor by residualizing FR_{it} against the same fixed effects, and verifies that the comovement is unchanged.

Specifically, we decompose

$$FR_{it} = \alpha_i + \kappa_{\text{tenure}(it)} + \tau_t + \widetilde{FR}_{it}, \quad (25)$$

and define the sign-based proxy on the residualized revision,

$$\widetilde{U}_{it} \equiv \mathbf{1}\{\widetilde{FR}_{it} > 0\},$$

following the same convention as in Section 2.2.2. The estimating equation is

$$\Delta Y_{it} = \alpha_i + \kappa_{\text{tenure}(it)} + \tau_t + \gamma_1 |\widetilde{FR}_{it}| + \gamma_2 (\widetilde{U}_{it} \cdot |\widetilde{FR}_{it}|) + u_{it}, \quad (26)$$

where α_i , $\kappa_{\text{tenure}(it)}$, and τ_t are household, panel-tenure, and month fixed effects.

Table 8 reports the estimates. Columns (1) and (2) use the change in the interquartile range and the change in the variance, respectively, as the dependent variable. In both columns, γ_1 is negative and $\gamma_1 + \gamma_2$ is positive, with both statistically significant: subjective uncertainty falls with downward (favorable) residualized revisions and rises with upward (unfavorable) ones, confirming the positive comovement predicted by Proposition 3.9. The estimates are quantitatively almost identical to those obtained with the raw revision in Section 2.3, consistent with the absence of substantial bias from reference-mean drift components in this sample.

Table 8: Point-forecast revisions and subjective uncertainty in the SCE: residualized specification

	(1)	(2)
	ΔIQR_{it}	ΔVAR_{it}
Favorable (γ_1)	-0.365*** (0.010)	-1.906*** (0.073)
Unfavorable ($\gamma_1 + \gamma_2$)	0.313*** (0.013)	1.547*** (0.094)
Household FE	Yes	Yes
Panel-tenure FE	Yes	Yes
Month FE	Yes	Yes
Observations	105,965	105,965

Notes: The dependent variable in column (1) is the consecutive-month change in the interquartile range of the household's year-ahead inflation density forecast, $\Delta IQR_{it} = IQR_{it} - IQR_{i,t-1}$. The dependent variable in column (2) is the corresponding change in the variance measure, $\Delta VAR_{it} = VAR_{it} - VAR_{i,t-1}$. The regressors are the absolute residualized revision $|\widehat{FR}_{it}|$ obtained from (25) and its interaction with the sign-based proxy $\tilde{U}_{it} = \mathbf{1}\{\widehat{FR}_{it} > 0\}$. The reported coefficients are the favorable-direction slope γ_1 and the implied unfavorable-direction slope $\gamma_1 + \gamma_2$, both obtained from the same regression. All specifications use SCE survey weights and absorb household, panel-tenure, and month fixed effects. Standard errors clustered two ways by household and month are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

References

- Afrouzi, Hassan, Spencer Yongwook Kwon, Augustin Landier, Yueran Ma, and David Thesmar. 2023. “Overreaction in Expectations: Evidence and Theory.” *The Quarterly Journal of Economics* 138(3): 1713–1764.
- Aktuğ, Emrehan, and Abolfazl Rezaghi. 2025. “Asymmetric Overreaction.” *European Economic Review* 180: 105153.
- Baqae, David Rezza. 2020. “Asymmetric Inflation Expectations, Downward Rigidity of Wages, and Asymmetric Business Cycles.” *Journal of Monetary Economics* 114: 174–193.
- Bhandari, Anmol, Jaroslav Borovička, and Paul Ho. 2025. “Survey Data and Subjective Beliefs in Business Cycle Models.” *The Review of Economic Studies* 92(3): 1375–1437.
- Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo. 2024a. “Diagnostic Business Cycles.” *The Review of Economic Studies* 91(1): 129–162.
- Bianchi, Francesco, Cosmin L. Ilut, and Hikaru Saijo. 2024b. “Smooth Diagnostic Expectations.” NBER Working Paper No. 32152.
- Binder, Carola, Pascal Frank, and Jane M. Ryngaert. 2025. “The Causal Effect of News on Inflation Expectations.” NBER Working Paper No. 34088.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2018. “Diagnostic Expectations and Credit Cycles.” *The Journal of Finance* 73(1): 199–227.
- Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer. 2019. “Diagnostic Expectations and Stock Returns.” *The Journal of Finance* 74(6): 2839–2874.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. “Overreaction in Macroeconomic Expectations.” *American Economic Review* 110(9): 2748–2782.
- Bordalo, Pedro, Nicola Gennaioli, Spencer Yongwook Kwon, and Andrei Shleifer. 2021. “Diagnostic Bubbles.” *Journal of Financial Economics* 141(3): 1060–1077.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2022. “Overreaction and Diagnostic Expectations in Macroeconomics.” *Journal of Economic Perspectives* 36(3): 223–244.
- Bordalo, Pedro, Nicola Gennaioli, Andrei Shleifer, and Stephen J. Terry. 2026. “Real Credit Cycles.” *American Economic Review* 116(4): 1274–1308.
- Carroll, Christopher D. 2003. “Macroeconomic Expectations of Households and Professional Forecasters.” *The Quarterly Journal of Economics* 118(1): 269–298.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. “Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments.” *American Economic Journal: Macroeconomics* 9(3): 1–35.

- Chahrour, Ryan, Adam Hale Shapiro, and Daniel J. Wilson. 2025. “News Selection and Household Inflation Expectations.” NBER Working Paper No. 33837.
- Coibion, Olivier, and Yuriy Gorodnichenko. 2015. “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts.” *American Economic Review* 105(8): 2644–2678.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2022. “Monetary Policy Communications and Their Effects on Household Inflation Expectations.” *Journal of Political Economy* 130(6): 1537–1584.
- Cornand, Camille, Paul Hubert, and Rose Portier. 2025. “Behavioral Underpinnings of Overreaction in Inflation Expectations Across Economic Agents.” HAL Working Paper hal-05133563.
- D’Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber. 2021. “Exposure to Grocery Prices and Inflation Expectations.” *Journal of Political Economy* 129(5): 1615–1639.
- D’Acunto, Francesco, Daniel Hoang, and Michael Weber. 2022. “Managing Households’ Expectations with Unconventional Policies.” *The Review of Financial Studies* 35(4): 1597–1642.
- D’Acunto, Francesco, Evangelos Charalambakis, Dimitris Georgarakos, Geoff Kenny, Justus Meyer, and Michael Weber. 2024. “Household Inflation Expectations: An Overview of Recent Insights for Monetary Policy.” NBER Working Paper No. 32488.
- Dräger, Lena, and Michael J. Lamla. 2017. “Imperfect Information and Consumer Inflation Expectations: Evidence from Microdata.” *Oxford Bulletin of Economics and Statistics* 79(6): 933–968.
- Duffy, John, Nobuyuki Hanaki, and Donghoon Yoo. 2025. “Overreaction in Expectations under Signal Extraction: Experimental Evidence.” Institute of Social and Economic Research, Osaka University Discussion Paper No. 1293.
- Ferland, Elyas, Camelia M. Kuhnen, Geng Li, and Itzhak Ben-David. 2024. “Extrapolative Uncertainty and Household Economic Behavior.” *Management Science* 70(8): 5607–5625.
- Hansen, Lars Peter, and Thomas J. Sargent. 2001. “Robust Control and Model Uncertainty.” *American Economic Review* 91(2): 60–66.
- Hansen, Lars Peter, and Thomas J. Sargent. 2008. *Robustness*. Princeton, NJ: Princeton University Press.
- Hansen, Lars Peter, Thomas J. Sargent, Gauhar Turmuhambetova, and Noah Williams. 2006. “Robust Control and Model Misspecification.” *Journal of Economic Theory* 128(1): 45–90.

- Huo, Zhen, Marcelo Pedroni, and Guangyu Pei. 2024. “Bias and Sensitivity under Ambiguity.” *American Economic Review* 114(12): 4091–4133.
- Ilut, Cosmin, and Hikaru Saijo. 2021. “Learning, Confidence, and Business Cycles.” *Journal of Monetary Economics* 117: 354–376.
- Ilut, Cosmin L., and Martin Schneider. 2014. “Ambiguous Business Cycles.” *American Economic Review* 104(8): 2368–2399.
- Kuchler, Theresa, and Basit Zafar. 2019. “Personal Experiences and Expectations about Aggregate Outcomes.” *The Journal of Finance* 74(5): 2491–2542.
- Lamla, Michael J., and Sarah M. Lein. 2014. “The Role of Media for Consumers’ Inflation Expectation Formation.” *Journal of Economic Behavior & Organization* 106: 62–77.
- Larsen, Vegard H., Leif Anders Thorsrud, and Julia Zhulanova. 2021. “News-Driven Inflation Expectations and Information Rigidities.” *Journal of Monetary Economics* 117: 507–520.
- L’Huillier, Jean-Paul, Sanjay R. Singh, and Donghoon Yoo. 2024. “Incorporating Diagnostic Expectations into the New Keynesian Framework.” *The Review of Economic Studies* 91(5): 3013–3046.
- Maxted, Peter. 2024. “A Macro-Finance Model with Sentiment.” *The Review of Economic Studies* 91(1): 438–475.
- Michelacci, Claudio, and Luigi Paciello. 2024. “Ambiguity Aversion and Heterogeneity in Households’ Beliefs.” *American Economic Journal: Macroeconomics* 16(2): 95–126.
- Pfajfar, Damjan, and Emiliano Santoro. 2013. “News on Inflation and the Epidemiology of Inflation Expectations.” *Journal of Money, Credit and Banking* 45(6): 1045–1067.
- Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent H. Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny, and Jorge Ponce. 2025. “Tell Me Something I Don’t Already Know: Learning in Low- and High-Inflation Settings.” *Econometrica* 93(1): 229–264.