FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

News Selection and Household Inflation Expectations

Ryan Chahrour
Cornell University

Adam Hale Shapiro and Daniel J. Wilson
Federal Reserve Bank of San Francisco

April 2025

Working Paper 2024-31

https://doi.org/10.24148/wp2024-31

Suggested citation:

Chahrour, Ryan, Adam Hale Shapiro, and Daniel J. Wilson. 2025. “News Selection and Household Inflation Expectations.” Federal Reserve Bank of San Francisco Working Paper 2024-31. https://doi.org/10.24148/wp2024-31

News Selection and Household Inflation Expectations*

Ryan Chahrour† Adam Hale Shapiro‡ Daniel J. Wilson§

April 6, 2025

Abstract

We examine how the media's systematic selection of reporting topics influences household responses to inflation news. In a model where households learn about inflation from news coverage, households account for news selection when forming their expectations. Because media are more likely to report on inflation when it is high, the model implies an asymmetric response to news: high-inflation news changes expectations more than low-inflation news. We test this implication using household panel data, and find that exposure to higher-prices news increases inflation expectations by 0.4 percentage point, while exposure to lower-prices news has no significant effect.

JEL classification: E3, D84

Keywords: Inflation Expectations, News Media, News Selection

*The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco, or the Board of Governors of the Federal Reserve System. We thank Kenza Daoudi, Gus Kmetz, Aditi Poduri and Carson Wolber for excellent research assistance. We thank Regis Barnichon, Olivier Coibion, Kristoffer Nimark, and Ricardo Reis for helpful discussions. The paper benefited from comments and suggestions by seminar participants at the Federal Reserve Bank of Dallas, the ASSA annual meetings, the SED annual meeting, and the National University of Singapore Macro Days conference.

†Cornell University, ryan.chahrour@cornell.edu

‡Federal Reserve Bank of San Francisco, adam.shapiro@sf.frb.org

§Federal Reserve Bank of San Francisco, daniel.wilson@sf.frb.org

1 Introduction

Inflation expectations are central to modern theories of inflation dynamics (Woodford, 2003). Empirical evidence shows that household inflation expectations, in particular, are a key driver of overall inflation (Coibion and Gorodnichenko, 2015), and are closely monitored by policymakers, particularly central banks. Household expectations also tend to be higher and more volatile than professional forecasts or market-based measures of expectations.

This study aims to provide new insight into how households form their inflation expectations by focusing on the role of the news media in shaping these expectations. The existing literature identifies several drivers of household inflation expectations. For example, households tend to overweight prices of salient, frequently purchased products such as food and energy (Trehan et al., 2011 and D'Acunto et al., 2021). Furthermore, personal experiences, including participation in the housing market (Kuchler and Zafar, 2019) and living during high inflation (Malmendier and Nagel, 2016), have been shown to impact household expectations. Although some prior research has considered the role of the media (e.g., Carroll, 2003 and Larsen et al., 2021), these studies generally rely on aggregate time-series data.[1]

A key concern in studying the effect of the media on expectations is the potential for reverse causality: the media may report more on inflation precisely when inflation or inflation expectations are elevated. Nimark and Pitschner (2019) formalize the systematic relationship between media reporting and economic conditions with the notion of a "news selection function." Systematic news selection not only makes it difficult to identify the effects of news coverage on expectations, it also potentially changes how those expectations respond to news: Because what is covered depends on the state of the world, the topics reported in the news are informative to households above and beyond the specific content of the stories themselves.

Motivated by the challenges of identifying the causal effects of media coverage on inflation expectations, we develop a model in which households learn about inflation through exposure to news. The model is calibrated to match the aggregate patterns of news coverage and household exposure seen in the data, and it delivers a clear testable prediction: favorable (i.e., lower prices) and unfavorable (i.e., higher prices) inflation news have asymmetric effects on household inflation expectations. We find support for this implication using micro panel data from the University of Michigan (UM) Survey of Consumers and find robust support.

Our model starts with the assumption that news outlets report more frequently on inflation when it is elevated, a pattern evident in the data.[2] In this framework, rational house-

[1]Ehrmann et al. (2017) also use household-level data, but do not include household or time fixed effects.

[2]Nimark (2014), Nimark and Pitschner (2019), and Chahrour et al. (2021) all focus on contexts where unusual events - whether high or low - are more likely to be reported.

holds encounter news stories drawn randomly from a distribution, in which some stories focus on inflation and others do not. As inflation increases, the probability that a randomly encountered story is about inflation increases. Consequently, the fact that a story addresses inflation conveys information beyond its specific content: hearing about inflation itself signals to a household that inflation may be higher than it previously thought. In addition to this inflation-*topic* signal, households also update their expectations based on the *content* of the story. The agent's inflation expectations decline if the story reports favorable conditions (i.e, lower prices); expectations rise if the story reports unfavorable conditions (i.e, higher prices).

As noted above, our model generates a distinctive prediction: the effects of favorable and unfavorable inflation news from the media on household inflation expectations are asymmetric. When households read unfavorable news, both the topic of the story and its content signal high inflation, leading to a substantial increase in expectations. By contrast, when households read favorable news, they receive conflicting signals: the story's focus on inflation suggests that inflation is a concern, but the content indicates low inflation. This leads to a more muted reduction—or no change at all—in inflation expectations.

To evaluate the importance of news-selection effects against the data, we include in our model an exogenous fraction of agents who deliberately seek inflation news. These households learn solely from the specific content—rather than the topic—of stories they observe. Unlike those who are randomly exposed to media coverage, inflation news seekers respond symmetrically to favorable and unfavorable inflation news. A calibration of the model demonstrates that substantial news-selection effects are necessary to match the aggregate time-series patterns in the data. Specifically, the model with only deliberate news seekers cannot simultaneously match two empirical patterns: the strong positive correlation between realized inflation and hearing unfavorable inflation news, and the weak negative correlation between realized inflation and hearing favorable news.

We test our model's household-level predictions, using the Michigan survey micro data to estimate the causal impact of hearing inflation news on expectations. The survey asks households whether they have heard news about business conditions and, if so, what they have heard. Households are free to respond that they have heard news about prices/inflation – the focus of our study – and, if so, whether the news was about *higher* (i.e., unfavorable) or *lower* (i.e., favorable) prices. The panel structure of the data, in which households are surveyed twice, six months apart, allows us to include both time and household fixed effects. These fixed effects help isolate the causal impact of the response to the survey question on inflation expectations. Time fixed effects control for the possibility that high aggregate inflation may simultaneously push up both household inflation expectations and news media coverage of inflation. Household fixed effects account for time-invariant household characteristics, such as demographic factors, that are correlated both with expectations and with the likelihood of hearing inflation news.

Key to our micro-level analysis is that households in a given survey round are asked the question about news twice. After asking the household about hearing news on business conditions (and what that news was about), the survey interviewer then asks the household if they have heard any *other* news. In the time-series data, we see that the share of households whose second response was "inflation news" (either favorable or unfavorable) is strongly correlated with the share of news stories mentioning inflation. By contrast, the share of households whose *first* response was "inflation news" is uncorrelated with news coverage of inflation. This result suggests that responses to the follow-up question are more sensitive to the *supply* of inflation coverage in the media, while responses to the primary question are more sensitive to idiosyncratic factors and/or individual's *demand* for inflation news (information-seeking).[3] The identifying assumption of our empirical approach is that individuals whose follow-up/secondary response was about inflation were not seeking out inflation news, but rather received it exogenously from the news media.

Our panel analysis under this identification strategy corroborates the model's predictions. In particular, we find symmetric effects on inflation expectations for households whose primary response was hearing inflation, but asymmetric effects for households whose secondary response was hearing inflation. One-year ahead inflation expectations for secondary responders rise by about 0.4 percentage point after hearing unfavorable inflation (higher prices) news and declines by a statistically insignificant 0.1 percentage point when hearing favorable inflation (lower prices) news. For primary responders (those whose first response to the news question was about inflation), one-year inflation expectations rise by about 0.2 percentage point after hearing unfavorable inflation news and declines by a statistically significant 0.3 percentage point when hearing favorable inflation news. Both the existence and sign of this bias on first responses is predicted by our model with a proportion of households that are randomly "inflation-curious."[4] Specifically, news selection acts to raise (i.e., move in a positive direction) the coefficients of the impact of favorable and unfavorable news on inflation

[3] This ordering of responses is intuitive, but also consistent with the psychology literature showing that individuals respond to open-ended questions with the topics that are most salient in their mind Zaller and Feldman (1992). The question ordering in the Michigan survey also addresses concerns about priming. The survey asks households about economic news *before* asking about their expectations. This reduces the risk of priming respondents to answer the economics news question in a way that rationalizes the expectations they have reported earlier in the survey.

[4] Inflation curiosity shocks (i.e., inflation-news seeking behavior) that are positively correlated with hearing unfavorable news (i.e., non-random curiosity shocks) would bias the coefficient upwards. However, our results on first responses show that that coefficient on unfavorable inflation news is lower in magnitude (i.e., absolute value) than that on hearing favorable inflation news.

expectations. The empirical results are robust to including time-varying household response controls (e.g., responses about personal finances and hearing news about non-inflation topics) and allowing demographic effects to vary by time-period (e.g., age effects may vary by time period). A simple back-of-the-envelope calculation suggests that media coverage of inflation could explain between 4 and 18% of the COVID-era increase in household inflation expectations.

Previous research has assessed whether and how news influences households' inflation expectations. In particular, Carroll (2003) finds evidence that the gap between household (Michigan) and professional (SPF) inflation expectations shrinks when news coverage of inflation is higher. He argues that inflation news informs households about professional forecasts, leading households to revise their expectations towards professionals. More recently, Larsen et al. (2021) shows that news media coverage of certain topics (such as technology and health care) are positively correlated with inflation expectations and negatively correlated with households' inflation forecast errors, suggesting that aggregate news may provide information that improves households' expectations. However, these studies rely on aggregate time-series, which makes establishing a causal relationship difficult.

A recent literature has also estimated the causal effects of information/news treatments on inflation expectations using randomized control trials (RCTs) (e.g., Armantier et al., 2016; Cavallo et al., 2017; Coibion et al., 2022; Weber et al., 2025 among many others). Though experiments can be very useful in assessing expectation responses to new information, the particular mechanism we study would be difficult to identify in a lab setting. Identifying the news selection mechanism in an experiment would require that survey respondents understand how the *information provider* selects what information to supply and when. An experiment looking for news selection effects would therefore need either to credibly communicate that the information treatment is randomly selected from media coverage at a given moment[5] or otherwise convince the survey participants that the probability of seeing inflation-related information is strongly increasing in current inflation.

The study is organized as follows. In section 2, we describe the UM data and motivating facts. Namely, that households are more likely to report hearing both unfavorable and favorable news about inflation when news media coverage of inflation is high. In section 3, we introduce the model of news selection and compare its predictions to the aggregate data. In sections 4 and 5, we describe our causal estimation regression and report the estimates. We conclude in section 6.

[5]In which case the sample size requirements would be huge, since only a small fraction of the sample would be exposed to inflation news.

2 Data and Motivating Facts

Our main data source is the University of Michigan (UM) Surveys of Consumers. In addition to the commonly-used question regarding inflation expectations, this survey includes a question about overall business conditions. Specifically, the survey asks: "During the last few months, have you heard of any favorable or unfavorable changes in business conditions?" The household can respond "Yes" or "No; Haven't Heard." If the respondent says "Yes," they are then asked a follow-up question regarding the type of news: "What did you hear? (Have you heard of any other favorable or unfavorable changes in business conditions?)." The Appendix contains an image of the survey questionnaire.

Table 1: Summary Statistics: Household Survey

	Full Sample		Heard unfavorable inflation news		Heard favorable inflation news	
	mean	sd	mean	sd	mean	sd
Expected Inflation, 1-year	4.62	3.99	5.86	4.54	3.56	3.42
Heard any news?	.653	.476				
(1st response)	.653	.476				
(2nd response)	.405	.491				
Heard unfavorable inflation news	.056	.23				
(1st response)	.0329	.178				
(2nd response)	.023	.15				
Heard favorable inflation news	.0145	.12				
(1st response)	.00832	.0908				
(2nd response)	.00623	.0787				
Male	.527	.499	.594	.491	.705	.456
Age	48	16.9	50.3	15.9	46.1	15.7
College Grad	.471	.499	.586	.493	.667	.471
HS Grad	.254	.435	.253	.435	.211	.408
Have Kids	.356	.479	.336	.472	.378	.485
Married	.632	.482	.665	.472	.677	.468
Income	66436	69150	91548	88662	77406	81451
Observations	138786		7768		2019	

The interviewer then codifies the responses by favorability (i.e., "favorable" or "unfavorable") and topic (e.g., government, employment, or inflation). The second part of the follow-up question, "Have you heard of any other...", allows the respondent to list an additional topic. We focus on those households who responded with news about inflation, which UM codes as "Unfavorable news: Higher prices" and "Favorable news: Lower prices."[6]

[6] We removed the very few observations where the individual reported hearing "Unfavorable news: lower

Figure 1: Aggregate variation in hearing inflation news

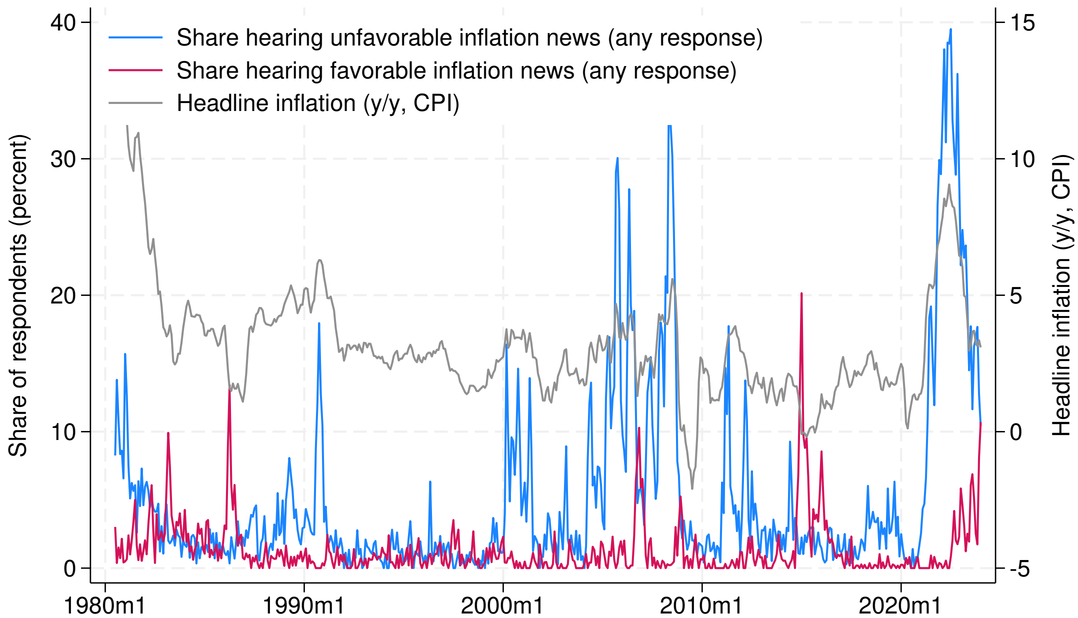

Table 1 shows summary statistics from the UM survey. Our analysis focuses on the panel nature of the survey and therefore isolates those households that are sampled at least twice, which amounts to 138,786 observations covering January 1980 to January 2024.[7] Approximately two-thirds (65.3%) of observations report hearing news about business conditions, of which a subset report hearing inflation news. Specifically, 7,768 observations (5.6%) report hearing unfavorable inflation news (i.e., news of higher prices) and 2,019 observations (1.5%) report hearing favorable inflation news (i.e., news of lower prices). Those households reporting unfavorable inflation news have higher-than-average one-year ahead inflation expectations, 5.9%, while those reporting favorable inflation news have lower-than-average inflation expectations, 3.6%. One-year ahead inflation expectations average 4.6% over the entire sample. Those reporting unfavorable news tend to be demographically similar to the average, while those reporting favorable news are more likely to be male and college graduates.

We also report the order in which inflation was reported by the respondent. Overall, 41% of observations include a second topic. Of the 5.6% of observations with unfavorable

prices" or "Favorable news: higher prices." No results changed when these observations were included. See appendix figure A3 for an image of the survey codebook.

[7]We removed outliers: the top and bottom 1 percent of households based on their inflation expectations level and the change in expectations over the panel period.

Table 2: Aggregate variation in hearing inflation news

	(1) Share Unfavorable	(2) Share Unfavorable	(3) Share Favorable	(4) Share Favorable
Core inflation (3-month, CPI)	0.484***	0.111	0.0650*	-0.0733*
	(0.106)	(0.122)	(0.0373)	(0.0435)
Food inflation (3-month, CPI)	0.945***	0.728***	-0.0371	-0.117***
	(0.103)	(0.0939)	(0.0276)	(0.0304)
Energy inflation (3-month, CPI)	0.0708***	0.0792***	-0.0393***	-0.0362***
	(0.0150)	(0.0142)	(0.00718)	(0.00700)
Unleaded gasoline price (per gallon)	3.907***	3.646***	0.186**	0.0890
	(0.299)	(0.259)	(0.0849)	(0.0781)
News coverage mentioning inflation (share)		1.551***		0.575***
		(0.269)		(0.0965)
N	523	523	523	523
R^2	0.580	0.615	0.152	0.211
Adjusted R^2	0.576	0.611	0.146	0.203

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is the share of households reporting heard news of higher prices (columns 1 and 2) and lower prices (columns 3 and 4), by month.

news a bit over half, 3.3 percentage points (pp), were reported as the first/primary response. Similarly, for favorable news, a little over half were reported as the first response. The distinction between those who mention inflation first and second will play an important role in our subsequent analysis.

As shown in Figure 1, the fraction reporting inflation news varies over time, with unfavorable news tending to rise with inflation and favorable news declining, however to a much smaller extent. Specifically, the correlation between the share of unfavorable news and year-over-year headline CPI inflation is 0.5, while the correlation with favorable news is just -0.05. The share of unfavorable news peaked in May 2022, at 43%. It also spiked, rising above 30%, in the summer of 2008 as inflation briefly climbed above 5%.

Table 2 provides an overview of the factors associated with time series fluctuations in the heard-news variables. Column (1) shows that the share of households that report hearing news of higher prices is positively associated with 3-month CPI inflation, separated into the core, food, and energy categories. The regression also includes the level of gas prices, as previous studies have shown that the inflation expectations are sensitive to this variable (Coibion and Gorodnichenko, 2015). Column (2) adds a variable measuring inflation news penetration taken from a corpus of news articles from 26 major newspapers across the United States (see Kmetz et al. (2022) for a description of the news media data). Specifically, this news coverage of inflation (NCI) variable represents the fraction of mentions of the term "inflation" across all articles. The coefficient on this variable is positive and significant, and incorporating it leads to an insignificant coefficient on core inflation. This pattern suggests that households learn about core inflation through the news media. Food and energy inflation and the price of gasoline, on the other hand, continue to be positively associated with the share hearing news about higher prices, consistent with the idea that information on these items is readily observable to consumers and hence they do not need the news media to learn about them.

Columns (3) and (4) report the same regressions but with the dependent variable being the share of households hearing favorable inflation news (i.e., news of lower prices). The first coefficient in Column (3) shows that the share hearing favorable inflation news *rises* slightly with core inflation. This relationship is much weaker than the one reported in Column (1), and the sign is somewhat counterintuitive: we might expect news about lower prices to fall when inflation is higher. In the next section, we show that our model can capture the relatively weak relationship displayed in this column. Column (4) shows that, as with case of hearing about unfavorable news, news media coverage soaks up the positive correlation between core inflation and hearing favorable inflation news. These results indicate that core inflation does not have an effect on people's reporting of hearing favorable inflation news independent of the news media channel. Overall, columns (1) and (3) show that households are more likely to report hearing both unfavorable and favorable news about inflation when inflation is high, while columns (2) and (4) show that this association can be accounted for by news media coverage.

While news coverage is strongly correlated with survey participants' reports of inflation news heard, Table 3 shows that this connection is much stronger for those who report they heard some other type of news first. The table shows the results of a regression of the NCI on measures of households' reported exposure to inflation news. The first column corroborates the result from Table 2 that the share of households reporting that they have heard inflation news (whether favorable or unfavorable) is strongly correlated with NCI. The second column shows that this relationship is both stronger and far more precise for households' second responses. The fact that the connection between news coverage and news exposure is much stronger for second responses motivates us to think about these responses separately in our theoretical analysis below.

Table 3: Aggregate variation in News Coverage Mentioning Inflation (NCI)

	(1)	(2)
Share hearing unfavorable inflation news (any response)	0.0751***	
	(0.0108)	
Share hearing favorable inflation news (any response)	0.119***	
	(0.0285)	
Share hearing unfavorable inflation news (1st response)		-0.00874
		(0.0200)
Share hearing favorable inflation news (1st response)		0.0385
		(0.0600)
Share hearing unfavorable inflation news (2nd response)		0.232***
		(0.0342)
Share hearing favorable inflation news (2nd response)		0.254***
		(0.0851)
N	523	523
R^2	0.247	0.303
Adjusted R^2	0.244	0.298

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is share of news stories mentioning inflation.

3 A Model of News Selection and Inflation Expectations

The time series evidence summarized above suggests that the supply of news about inflation varies over time, rising in periods when inflation itself is above average. Table 3 also suggests that people report reading news about inflation for more than one reason. In this section, we provide a simple model of time variation in inflation coverage, and explore how fluctuations in the intensity of inflation coverage impact the inferences people draw from the news. The key ingredient of the model is a "news selection function" in which media are more likely to report on inflation when inflation is high.

To study the model's implications, we emphasize two special cases of the model. In the first, media news selection—even if present—has no effect on the types of news stories seen by households. In the second, we allow the selection of news reporting to influence the stories seen by households. Because households understand that the media select the topics they report based on the state of the world, the topic of a story they see provides information beyond that which is explicitly contained in the story itself. We conclude the section by examining a version of the model that nests both special cases.

3.1 Model Environment

Households are indexed by $i \in \mathcal{I}$. Each period, households are asked to forecast the end-of-period realization of inflation, π_{t+1}. Household i's inflation expectation is defined as

$$\pi^e_{it} \equiv E[\pi_{t+1}|\Omega_{it}]. \tag{1}$$

Inflation is an i.i.d. random variable with symmetric probability distribution function $f(\pi_{t+1})$ and unconditional mean $\bar{\pi} \equiv E[\pi_{t+1}]$.[8] The information set of household i is denoted Ω_{it}, which we define below.

Each period, the media generate a continuum of news stories indexed by $m \in \mathcal{M}_t$. Some stories are about inflation while others concern different topics: the set of stories that are about inflation is denoted by $\Pi_t \subseteq \mathcal{M}_t$. We assume that the topics covered in the media are determined by a "news selection function" (Nimark and Pitschner, 2019). In particular, the probability of a story m being about inflation is given by

$$Pr[m \in \Pi_t|\pi_{t+1}] = \Phi(\pi_{t+1}). \tag{2}$$

The function $\Phi(\cdot)$ is called the "news selection function" and it describes how inflation changes the likelihood that a given media story is about inflation.[9] We assume throughout that $\Phi(\pi_{t+1})$ is a strictly increasing function.

Every report on the topic of inflation contains an unbiased signal about inflation over the period. The signal about inflation contained in an inflation story is given by

$$s_{mt} = \pi_{t+1} + \eta_{mt}, \tag{3}$$

where the noise $\eta_{mt} \sim N(0, \sigma^2_\eta)$ is independent across households. We assume that stories that are not about inflation contain no economic information at all.

To map the model to the data, we define "favorable news" and "unfavorable news" about inflation. We assume that a signal constitutes favorable news if it suggests inflation will be below average

$$s_{mt} < \bar{\pi}$$

and a signal is "unfavorable news" if it suggests inflation will be above average.[10]

[8] Given the static nature of our model, we use the notions of "inflation" and "higher prices" interchangeably.

[9] We model the topic of a particular report as a probabilistic, and not a *deterministic*, function of inflation because it is assumed that other newsworthy events (i.e., events unrelated to inflation) occur randomly, as discussed in Nimark and Pitschner (2019).

[10] Our association of higher inflation with unfavorable news contrasts with the U-shaped loss function

In a period, a household hears at most two of the news stories in $\mathcal{M}_t$. The choice to allow households to hear up to two stories is motivated by the structure of the Michigan survey, in which households are permitted to report hearing up to two categories of economics news.

The topics of the first and second stories observed are selected using different mechanisms. In particular, we assume that the topic of the first story is selected based on the households' desired to read about it: the first topic is determined by the households' "demand" for that information. Meanwhile, we assume that the topic of the second story is determined by households' random encounters with news stories: thus, the second topic is determined by the "supply" of information in the media. Distinguishing mechanisms between the first and second stories is pedagogically useful in highlighting the importance of news selection effects in the data. Our assumptions about the response order also align with the psychological evidence that households first respond to open-ended questions with the topics that are most salient in their minds (Zaller and Feldman, 1992)—that is, news the household intentionally pursued. The assumption that first responses are driven by households' demand for the inflation topic and not the supply of inflation stories is supported by the evidence in Table 3.

All households observe a first news story. The topic of that story is determined by a random shock to the household's preferences across topics. Specifically, with probability ϕ, the household's preferred topic is inflation. In this case, we say the household received an "inflation curiosity shock". With probability $1 - \phi$, the household is curious about another topic and seeks out information unrelated to inflation. Inflation-curious households knowingly consume an inflation news story indexed by $m_a(i) \in \Pi_t$, which is selected randomly from the set of all inflation stories. Observing the signal $s_{m_a(i),t}$, the household updates their inflation expectations optimally. The set of households that *intentionally* observe a story about inflation is denoted $\mathcal{A}_t$. A household not experiencing an inflation-curiosity shock intentionally observes a non-inflation news story, indexed by $m_a(i) \notin \Pi_t$. Since non-inflation stories contain no useful economic information, the household does not update their priors about inflation at all.

Households may receive a second news story with exogenous probability ω. If a second story is received, the topic is determined by the relative supply of different topics in the media. In particular, if the household has not already intentionally consumed inflation news, then the household observes a random news story $m_b(i) \in \mathcal{M}_t$. When the household observes its

around a positive inflation target that describes most policy makers' preferences (Shapiro and Wilson, 2022). However, this convention is consistent with the fact that households in the Michigan survey almost never report "Unfavorable news: lower prices" or "Favorable news: higher prices." It is also consistent with Afrouzi et al. (2024)'s finding that surveyed households prefer an inflation rate near zero (0.2%), implying lower (higher) prices are always favorable (unfavorable) so long as inflation is positive.

story, it learns whether the story is about inflation or not, i.e., whether $m_b(i) \in \Pi_t$. If it is, then they also see the signal $s_{m_b(i),t}$. Because the second story is randomly selected from the population of news stories, the probability that $m_b(i) \in \Pi_t$ is equal to the aggregate share $\phi(\pi_{t+1})$ of current stories that are about inflation. The set of households that *unintentionally* hear a story about inflation is denoted $\mathcal{B}_t$.

Because of news selection, *observed second stories are more likely to be about inflation whenever inflation is higher.* Households whose randomly-selected news story is about inflation therefore make inference from the topic of the coverage, as well as the signal it reports. This is the key difference relative to standard inference problems, including the inference problem associated with intentionally-chosen first stories: the systematic response of news media reporting to economic conditions means that people learn from the *topic* of the news they receive, above and beyond the particular *content* of the story.

3.2 Model Implications

To elucidate the implications of this information structure for household inflation expectations, we focus first on two special cases of the model. In the first case, we eliminate the effects of news selection by assuming that the probability that households receive a second news story, ω, is zero: in this case, only households that specifically seek out inflation news observe it. In the second special case, we isolate the effects of news media's selective reporting by setting the share of "inflation curious" households, ϕ, to zero. In this case, only the supply of inflation news determines the household's probability of observing inflation stories.

We conclude the section by studying the more general version of the model that includes households learning via both channels. We show that that this model implies a reduced-form representation with a particular interaction effect that does not appear in either of the special cases.

Model without news selection effects

In this version of the model, we set $\omega = 0$. In this case, only households experiencing inflation-curiosity shocks observe inflation news stories. Households know if they are inflation curious, and understand that the relative supply of inflation news stories does not impact their likelihood of reading about the inflation topic. In this sense, this version of the model assumes that inflation news consumption is completely "demand determined."

Because households in this version of the model choose to read a story about inflation, the topic of their news consumption has no impact on their inference. As a result, the household processes its inflation signal in a standard way: the household responds to inflation signals in a symmetric way, increasing their inflation expectations whenever $s_{m_a(i)} > \bar{\pi}$ and decreasing them otherwise. The first proposition summarizes this result:

Proposition 1. *In the model with $\phi > 0$, $\omega = 0$, a household reporting "unfavorable news" increases its expectations by the same amount as a household reporting "favorable news" decreases theirs. That is*

$$E[\pi^e_{it}|i \in \mathcal{A}_t, s_{m_a(i),t} > \bar{\pi}] - \bar{\pi} = -\left(E[\pi^e_{it}|i \in \mathcal{A}_t, s_{m_a(i),t} < \bar{\pi}] - \bar{\pi}\right). \tag{4}$$

Proof. Proved in the appendix. □

The argument for Proposition 1 is especially straightforward when shocks are normally distributed. A household that does not receive a curiosity shock has expectation equal to the unconditional mean, $\bar{\pi}$. Households that do receive a curiosity shock solve a standard linear-Gaussian filtering problem, with expectations given by $\pi^e_{it} = \bar{\pi} + \kappa(s_{m_a(i),t} - \bar{\pi})$ and κ a function of the signal-to-noise ratio of the story content. We can then compute

$$\begin{aligned} E[\pi^e_{it}|i \in \mathcal{A}_t, s_{m_a(i),t} > \bar{\pi}] - \bar{\pi} = & \ \ \kappa\left(E[s_{m_a(i),t} - \bar{\pi}|i \in \mathcal{A}_t, s_{m_a(i),t} > \bar{\pi}]\right) \\ = & -\kappa\left(E[s_{m_a(i),t} - \bar{\pi}|i \in \mathcal{A}_t, s_{m_a(i),t} < \bar{\pi}]\right) \\ = & -\left(E[\pi^e_{it}|i \in \mathcal{A}_t, s_{m_a(i),t} < \bar{\pi}] - \bar{\pi}\right) \end{aligned}$$

where the middle equality follows from the symmetry of the distribution of the signal $s_{m_a(i),t}$ around $\bar{\pi}$. We provide a more general proof in the appendix.

Using Proposition 1 and defining variables appropriately, we can derive a simple reduced-form expression describing model expectations depending on whether households have heard favorable or unfavorable inflation news. Define U^a_{it} as an indicator variable equal to one for those households i who have received an unfavorable signal about inflation $s_{m_a(i),t} > \bar{\pi}$. Let F^a_{it} be an indicator for those who have received favorable inflation news. Corollary 1 summarizes the resulting reduced-form representation. A detailed derivation appears in the Appendix.

Corollary 1. *When $\phi > 0$ and $\omega = 0$, the model implies a reduced-form representation of the form*

$$\pi^e_{it} = \alpha^a + \beta^a U^a_{it} + \gamma^a F^a_{it} + \mu_{it}. \tag{5}$$

where $\beta^a = -\gamma^a$ and μ_{it} is a mean-zero error term uncorrelated with the regressors.

The coefficients in (5) have clear structural interpretations as conditional averages of household forecasts. Specifically, in this case, the constant term is $\alpha^a = \bar{\pi} = E[\pi^e_{it}|i \notin \mathcal{A}_t]$ and the slope coefficients are given by $\beta^a = E[\pi^e_{it} - \bar{\pi}|i \in \mathcal{A}_t, s_{m_a(i),t} > \bar{\pi}]$, and $\gamma^a = E[\pi^e_{it} - \bar{\pi}|i \in \mathcal{A}_t, s_{m_a(i),t} < \bar{\pi}]$.

Model with news selection effects

We now consider a version of the model in which $\phi = 0$. In this case, the household only observes inflation news stories when it randomly draws such stories from the population of all news stories. Because the probability of drawing an inflation news story depends on the degree of news coverage, and news coverage depends on the current level of inflation, households learn something when they learn whether (or not) the story they have drawn concerns inflation.

To see how this works, it is useful to compare the (average) expectations of households that observe a non-inflation story to those who see a story about inflation, prior to learning anything about the signals contained inside the stories they see. Proposition 2 characterizes these average expectations:[11]

Proposition 2. *In the model with $\phi = 0$ and $\omega > 0$, the average household that observes a non-inflation story expects inflation to be below average, while the average household that observes an inflation story expects inflation to be above average. That is,*

$$E[\pi^e_{it}|i \in \mathcal{B}_t] > \bar{\pi} > E[\pi^e_{it}|i \notin \mathcal{B}_t]. \tag{6}$$

Proof. Proved in the appendix. □

The proof of Proposition 2 in the appendix makes clear that the proposition also describes what households themselves learn *just from learning the topic of a news story.* Prior to seeing the realization of its signal, a household that sees that the topic of a story is inflation should raise its inflation expectations above the prior (average) level of inflation, while the household that sees an unrelated story should lower its expectations below the average. Intuitively, the positive slope of the news selection function implies that the conditional distribution $f(\pi_{t+1}|i \in \mathcal{B}_t)$ always puts strictly more weight on higher values of π_{t+1}, compared to the unconditional distribution. Conversely, $f(\pi_{t+1}|i \notin \mathcal{B}_t)$ shifts weight to lower values of inflation.

To fully evaluate the rational expectations of a household that sees an inflation story, we must combine its pessimistic "prior" based on Proposition 2 with the information contained

[11] A similar result appears as Proposition 5 in Chahrour et al. (2021).

in the story, $s_{m_b(i),t}$, itself. Because our data provide only a qualitative classification of news as "favorable" or "unfavorable", we again focus on computing the average expectations depending on whether the signal is below or above average inflation. The difference of favorable news or unfavorable news households, relative to the "no news" households, will be one of the key outcomes in our empirical exercises. Proposition 3 provides an general characterization of the effects of seeing favorable or unfavorable news.

Proposition 3. *In the model with $\phi = 0$ and $\omega > 0$, households hearing unfavorable news increase their inflation expectation by more than households hearing favorable news decreases theirs, relative to "no news" households. That is,*

$$|E[\pi_{it}^e | i \in \mathcal{B}_t, s_{m_b(i),t} > \bar{\pi}] - E[\pi_{it}^e | i \notin \mathcal{B}_t]| > |E[\pi_{it}^e | i \in \mathcal{B}_t, s_{m_b(i),t} < \bar{\pi}] - E[\pi_{it}^e | i \notin \mathcal{B}_t]| \quad (7)$$

Proposition 3 is proved in the Appendix. Intuitively, if the agent sees "favorable news", i.e., news about lower inflation, then the content of the signal works against the signal sent by the inflation topic: the fact that the story is about inflation increases the household expectation of inflation, while the signal content itself indicates lowers inflation.

By contrast, the agent seeing "unfavorable news" experiences two forces that both increase inflation expectations: the fact that the news is about inflation, and the fact that signal itself indicates higher inflation. The contrast, between opposing forces in the first case, and reinforcing forces in the second, leads to the asymmetry in the response of households to seeing favorable versus unfavorable inflation news.

As before, our theory implies a simple reduced form relationship between the type of news observed by households and the their inflation expectations. Again defining U_{it}^b as a dummy variable for household hearing unfavorable news (i.e., $i \in \mathcal{B}$ and $s_{m_b(i),t} > \bar{\pi}$) and F_{it}^b as a dummy for those hearing favorable news, we have

Corollary 2. *When $\phi = 0$ and $\omega > 0$, the model implies a reduced-form representation of the form*

$$\pi_{it}^e = \alpha^b + \beta^b U_{it}^b + \gamma^b F_{it}^b + \mu_{it}. \quad (8)$$

where $\beta^b > 0$, $\gamma^b < 0$, $|\beta^b| > |\gamma^b|$, and μ_{it} is an mean-zero error term uncorrelated with the regressors.

In this case, the reduced-form coefficients correspond to different conditional expectations, namely $\alpha^b = E[\pi_{it}^e | i \notin \mathcal{B}_t]$, $\beta^b = E[\pi_{it}^e | i \in \mathcal{B}_t, s_{m_b(i),t} > \bar{\pi}] - \alpha^b$ and $\gamma^b = E[\pi_{it}^e | i \in \mathcal{B}_t, s_{m_b(i),t} < \bar{\pi}] - \alpha^b$.

The coefficients in (8) differ from those described in the reduced-form representation (5) in two important ways. First, in (8) the constant α^b is always lower than $\bar{\pi}$, reflecting that households who see non-inflation news infer that inflation is below average. Second, the effects of unfavorable and favorable news in (8) are asymmetric, for the reasons described above.

Full Model

Our baseline model, and the version that best captures the forces we identify in the data, includes a mix of inflation-seeking behavior (captured by the topic of the first story observed) and exogenous exposure to inflation news driven by the degree of inflation news coverage (captured by the second story observed.) Households that learn from inflation news they have specifically sought out respond symmetrically to favorable and unfavorably inflation news, while households that receive inflation news exogenously, but with a probability that increases with actual inflation, respond differently to favorable and unfavorable inflation news.

We can use the extended model to derive a "structural equation" describing household inflation expectations depending on the type of news they have observed. Let c_{it} be a dummy variable equal to one if the agent receives a curiosity shock, and therefore actively seeks out inflation news. Further, define the dummy U_{it} equal to one if the agent observed unfavorable inflation news ($s_{it} > 0$) in either the first (selected) or second (exogenously received) news story and define F_{it} to be the analogous variable for favorable inflation news.

Since this model contains a mix of households learning about inflation via two distinct channels, we derive a reduce-form expression for households' expectations by combining the reduced-form equations in (5) and (8). The resulting equation is given by the following corollary

Corollary 3. *When $\phi > 0$ and $\omega > 0$, the model implies a reduced-form representation of the form*

$$\pi_{it}^e = \underbrace{c_{it}[\alpha^a + \beta^a U_{it} + \gamma^a F_{it}]}_{\text{received curiosity shock}} + \underbrace{(1 - c_{it})[\alpha^b + \beta^b U_{it} + \gamma^b F_{it}]}_{\text{didn't receive curiosity shock}} + \mu_{it}, \tag{9}$$

where μ_{it} is a mean-zero error term uncorrelated with the regressors.

Equation (9) provides a nearly operational estimation equation. In expression (9), the coefficients continue to correspond to the same conditional expectations that appear in (5) and (8). Notice, as well, that premultiplication by the dummy c_{it} means that in (9) it is innocuous to replace U_{it}^a with U_{it}, and so on for all of the news heard dummy variables. According to the equation, if we estimate separate responses to good news and bad news reported in the first or second answers, we should expect symmetric responses between good and bad news if the respondent answers with inflation first, and asymmetric responses to second answers.

Collecting the terms multiplied by c_{it}, equation (9) can be rewritten as

$$\pi_{it}^e = \alpha^b + \underbrace{c_{it}[(\beta^a - \beta^b + d\alpha)U_{it} + (\gamma^a - \gamma^b + d\alpha)F_{it}]}_{ov_{it}} + \beta^b U_{it} + \gamma^b F_{it} + \mu_{it}, \tag{10}$$

where $d\alpha \equiv \alpha^a - \alpha^b$. Equation (10) is useful because it highlights why pooling responses, and estimating just one coefficient on e.g. unfavorable news, is perilous. Evidently, such an estimation would not deliver consistent estimates of β^b or γ^b because such a pooled regression would omit the interaction term, captured by ov_{it}. In section 4, we analytically derive the direction of this bias implied by our model.

3.3 A Numerical Example

To make the connections between our theory and our empirical findings more concrete, we report some implications from a numerical example of our model. Our general approach is to calibrate the model to match patterns from the aggregate time series data, and then explore the implications of the model for the individual-level estimation that we undertake using the micro data. The model is too simple to provide a realistic quantitative match to the inference problem being solved by households in practice, but it provides important qualitative direction for our empirical exercises.

For the process of inflation, we assume that $\pi_{t+1} \sim N(\mu_\pi, \sigma_\pi^2)$, with a mean $\mu_\pi = 2.85$ and standard deviation $\sigma_\pi = 1.60$ that are equal to those of realized CPI inflation in our sample. The remaining parameters concern the processes for media coverage and information acquisition. We calibrate these parameters to best match the time series relationship between the level of inflation and the share of household observing favorable or unfavorable inflation news. The raw data (and the corresponding bin-scatter) are reflected in the blue x's (and blue dots) in Figure 2.

To emphasize the important role of news selection in matching the time series data, we consider two calibrations. The first calibration corresponds to the special case of the model considered in section 3.2 where we assume that all households observing inflation news do so because they have received an inflation curiosity shock. In this case, we fix $\phi = 0.070$ to match the unconditional share of households who report hearing inflation news. We need only calibrate a value for the standard deviation of the idiosyncratic signal noise, σ_η. The

Figure 2: News selection in the model versus data

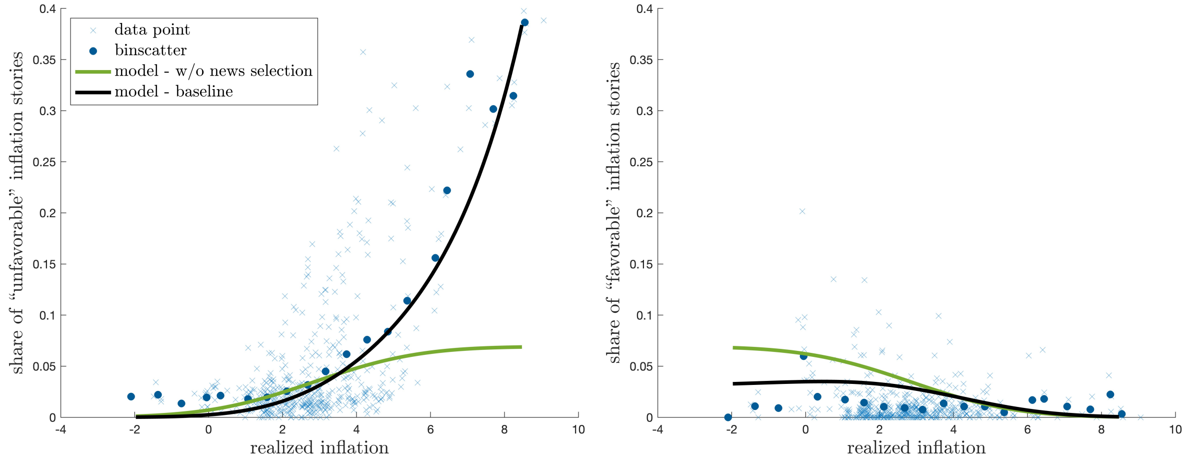

(a) Unfavorable news about inflation. (b) Favorable news about inflation.

Notes: Actual inflation data are year-over-year headline CPI inflation measured monthly from January 1983 to January 2024. Share of favorable/unfavorable news heard comes from the University of Michigan Survey of Consumers. Model share are based on analytical calculations from the calibrated distributions and news selection function.

values of this parameter is $\sigma_\eta = 2.270$.[12]

The solid green line in Figure 2 reflects the implied relationship between news heard and realized inflation for this calibration. The Figure shows that this version of the model does a rather poor job of accounting for the time series patterns of the two news-heard variables. In particular, while the share of unfavorable news heard is upwards sloping in both the data and the model, the model cannot match the extremely high exposure to inflation news during high inflation periods: the model-implied relationship is not upward-sloping enough. By contrast, for this version of the model, the share of favorable inflation stories appears to be too strongly downward sloping. When inflation is low, the model implies that households hear a relatively large share of favorable inflation stories, but the data reveal a modest increase, if any, of positive inflation news exposure during low inflation episodes. In short, the model without news selection effects cannot simultaneously match a strongly upward sloping relationship between inflation and unfavorable news heard and relatively flat relationship between inflation and favorable news heard.

Allowing for a news selection effect improves the fit of these time series properties sub-

[12] We select this parameters by minimizing the squared deviations between the data and model implied probabilities of observing unfavorable and favorable inflation news.

stantially. For our second (a.k.a. "baseline") calibration, we assume that on average 50% of households seeing inflation news do so because of curiosity shocks ($\phi = 0.035$); all remaining households experience time varying inflation news exposure according to the news selection function ($\omega = 1$). To calibrate the news selection function we use an exponential function[13]

$$\Phi(\pi_{t+1}) = ce^{d\pi_{t+1}}.$$

The parameter c targets the average fraction of news coverage about inflation, and d the degree to which coverage grows when inflation is high. We find that $c = 0.0246$ and $d = 0.4781$, along with a value of $\sigma_\eta = 1.949$, best match the time series data.

The black solid line in Figure 2 plots the implied relationship between realized inflation and the reports of unfavorable and favorable news in the correspond calibrated model. Panel (a) shows that the model does a much better job of matching the strong upward slope of unfavorable news in response to inflation. Moreover, Panel (b) shows that the fall in coverage at lower inflation levels mostly offsets the increasing share of inflation stories that are favorable, so that the share of favorable news heard is relatively flat (indeed, slightly non-monotonic) in realized inflation.

The model with a strong news selection effect is a better match for these features of the time series data. At the same time, Proposition 3 describes a clear cross-sectional implication of the model with strong news selection: the expectational impact of the *random observation of an inflation story* should be larger when that story suggests rising inflation. Analogously, responses to inflation stories that agents specifically sought out should be symmetric. In the next sections, we test these implications using the micro data from the Michigan Survey of Consumer Expectations.

4 Empirical Specification

Our empirical objective is to see if there is evidence of substantial news selection effects in how households update their expectations. To identify the effects of exposure to news, we turn to household-level panel data that includes information on both households' inflation expectations and whether they report hearing inflation news. Specifically, we construct a household panel data set using the microdata from the University of Michigan's Survey of Consumer Expectations (MSC). A panel is obtainable because a large number of households are sampled twice, six months apart.

[13]This specification is not bounded above by one, as a probability should be, but this constraint never binds in our sample or simulations.

Our use of household-level panel data offers three key identification advantages over an approach based purely on time series data, or one based purely on cross-sectional data. First, it allows us to absorb time fixed effects. These time fixed effects capture the influence of aggregate inflation and aggregate inflation expectations on inflation-related news coverage.[14] Second, we can utilize the individual-level measures of news described above that are available in the Michigan survey. Third, the micro panel data allows us to deal with a potential household-specific bias concern. Specifically, individuals that tend to have high (low) levels of inflation expectations may also tend to seek out inflation news more (less) so than others and, in turn, be more likely to report hearing news of higher prices (and less likely to report hearing news of lower prices). To the extent that such individual factors are fixed over the six months households are in the sample, they will be captured by individual fixed effects.

For estimation, we consider variants of the following regression specification:

$$\pi^e_{it} = \beta^a U^a_{it} + \gamma^a F^a_{it} + \beta^b U^b_{it} + \gamma^a F^b_{it} + \alpha_i + \alpha_t + \delta X_i \cdot t + \zeta Y_{i,t} + \epsilon_{it}, \tag{11}$$

where π^e_{it} is household i's self-reported one-year ahead inflation forecast as of month t. U^a_{it} is an indicator variable equal to 1 if the household reports having heard unfavorable inflation news (i.e., news of higher prices) in their first/primary response to the news survey question. Otherwise, U^a_{it} is set to zero. F^a_{it} is an analogous variable for favorable inflation news (i.e., hearing news of lower prices), and U^b_{it} and F^b_{it} are corresponding dummies for the secondary responses.

Relative to the structural equation (9), equation (11) allows for a variety of potentially important control variables. First, as noted above, we include a broad set of additional controls that eliminate individual- and time-specific differences that might influence inflation expectations for the reasons discussed above. In (11), the parameters α_i and α_t denote individual and time fixed effects, respectively. Although household fixed effects, α_i, control for demographics, we allow for the possibility that the effects of demographics on inflation expectations may vary over the six months separating households survey responses by including an interaction between each demographic variable and time, $X_i \cdot t$. Lastly, in some specifications, we include time-varying household-level variables, $Y_{i,t}$. These include responses to questions about personal finances and about hearing non-inflation news.[15]

[14] Figure A1 documents this correlation.

[15] The personal finance questions ask the respondent whether their personal finances are better or worse than a year ago, and whether they expect their personal finances to be better or worse a year from now. We also include controls for whether or not the household mentioned specific non-inflation related news in their response to the news question: hearing of plants opening or closing, hearing the stock market is rising or declining, hearing improvements or declines in specific industries, hearing that credit is easy to get, hearing that interest rates are high, and hearing that employment is high or low.

To increase efficiency, in practice we estimate (11) in first-differenced form:

$$\Delta\pi^e_{it} = \beta^a \Delta U^a_{it} + \gamma^a \Delta F^a_{it} + \beta^b \Delta U^b_{it} + \gamma^b \Delta F^b_{it} + \tilde{\alpha}_t + \delta X_i + \zeta \Delta Y_{i,t} + \varepsilon_{it}, \qquad (12)$$

where $\tilde{\alpha}_t = \Delta\alpha_t$. The difference operator, Δ, in equation (12) refers to a six-month difference because households in the Michigan survey are surveyed six months apart (if they are surveyed more than once at all).

The identifying assumption of our empirical approach is that individuals whose secondary response to the survey question was about inflation were not seeking out inflation news, but rather were exogenously supplied it by the news media. By contrast, households whose primary response to the news question is about inflation may either by supplied it exogenously or may have a specific demand for inflation news. The notion that households specifically interested in inflation news would answer about inflation news first is both intuitive and consistent with psychological literature on survey behaviors (e.g., Zaller and Feldman, 1992). It is also consistent with the time series evidence, presented in section 2, that news media coverage is only correlated with second responses to the news heard question. In section 5, we discuss why the qualitative patterns of our estimated coefficients provide additional support for our assumptions.

To the extent that some second responders actively seek out inflation news, our identification assumption will be violated. In this case, our model implies that the asymmetry we measure is a lower-bound on the degree of asymmetry caused by unintentional exposure to inflation news. To see this most clearly, consider the case where we have *no* information on whether households answered the question first or second. In that case, we would estimate the following pooled regression:

$$\pi^e_{it} = \beta U_{it} + \gamma F_{it} + \alpha_i + \alpha_t + \delta X_i \cdot t + \zeta Y_{i,t} + \epsilon_{it}, \qquad (13)$$

where $U_{it} = U^a_{it} + U^b_{it}$ and $F_{it} = F^a_{it} + F^b_{it}$.

In the appendix, we show that estimating (13) on data generated by the structural system in (10) delivers (population) estimates

$$\hat{\beta} = \beta^b + \frac{\phi}{2}\frac{\beta^a - \beta^b + d\alpha}{\mu^U} \qquad (14)$$

$$\hat{\gamma} = \gamma^b + \frac{\phi}{2}\frac{\gamma^a - \gamma^b + d\alpha}{\mu^F} \qquad (15)$$

where ϕ is the fraction of households receiving curiosity shocks, μ^U is the average fraction of people who observe unfavorable news and μ^F is the average fraction of people observing favorable news. According to our baseline model, $\beta^a - \beta^b - d\alpha < 0$ and $\gamma^a - \gamma^b - d\alpha < 0$. This implies the omitted variable term $c_{it}[(\beta^a - \beta^b)U_{it} + (\gamma^a - \gamma^b)F_{it}]$ would be negatively correlated with both U_{it} and F_{it}, causing a negative bias on both β and γ relative to β^b and γ^b. Intuitively, the estimates of the pooled regression reflect a mixture of the relatively symmetric effects of intentional/demand-driven news exposure and the asymmetric effects of exogenous/supply-driven exposure to inflation news. We report a version of this pooled regression as a reference point in the next section.

5 Empirical Results

5.1 Effects of News on Inflation Expectations

As a benchmark, we first present results of the pooled regression, shown in Table 4. Column (1) corresponds to the simplest version of the regression, where demographic controls (X_i), time-varying household controls ($Y_{i,t}$) and time fixed effects are omitted. Hearing unfavorable inflation news (i.e., news of higher prices) increases inflation expectation by about 0.6 percentage points. This represents about a 15 percent increase of inflation expectations relative to the sample mean of around 4.6%. The estimated coefficient on favorable news (i.e., hearing lower prices) inflation news is negative, approximately -0.4, indicating that hearing news of lower prices reduces inflation expectations. The coefficients on both favorable and unfavorable news are statistically significantly different from zero at below the 1 percent level. The inclusion of controls, shown in columns (2) through (7), attenuates the size of both coefficients (i.e., pushes the magnitudes of the coefficients on both unfavorable and favorable news towards zero). The p-value on the test of symmetry ($\hat{\beta} = -\hat{\gamma}$) is shown at the bottom of the table. Across all specifications, one cannot reject symmetry.

Table 5 shows results of estimating equation (12), which separates the treatment between first and second responses. Motivate by our model and time series observations, we expect that first and second responders should react differently to inflation news, with a difference that reflects the effects of news selection. Column (1) shows that a household that reports hearing unfavorable inflation news as their first response increases inflation expectations by about 0.5 percentage points. Hearing favorable news decreases its inflation expectations by about 0.5 percentage points—a symmetric effect of hearing favorable and unfavorable news on inflation expectations. Similar to the pooled estimates, including additional controls tends to attenuate the size of the coefficients but does not change the finding of symmetry. The specification with full controls, column (7), shows that hearing unfavorable inflation news increases inflation expectations by approximately 0.2 percentage points while hearing

Table 4: Pooled Treatment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Heard unfavorable inflation news	0.573***	0.428***	0.414***	0.291***	0.282***	0.273***	0.265***
	(0.0754)	(0.0692)	(0.0692)	(0.0654)	(0.0661)	(0.0657)	(0.0656)
Heard favorable inflation news	-0.406***	-0.318***	-0.319***	-0.227**	-0.249**	-0.245**	-0.236**
	(0.100)	(0.0975)	(0.0981)	(0.0974)	(0.0981)	(0.0981)	(0.0973)
Headline inflation (y/y)		0.269***					
		(0.0280)					
Core Inflation (y/y)			0.110**				
			(0.0450)				
Food inflation (y/y)			0.116***				
			(0.0227)				
Energy inflation (y/y)			0.0244***				
			(0.00342)				
Observations	67818	67818	67818	67818	65650	65650	65650
Adjusted R^2	0.002	0.009	0.010	0.000	0.002	0.003	0.004
$\lvert$unfav$\rvert$ = $\lvert$fav$\rvert$, p-val	0.184	0.368	0.437	0.586	0.777	0.811	0.803
Time Fixed Effects	No	No	No	Yes	Yes	Yes	Yes
Demographic Controls	No	No	No	No	Yes	Yes	Yes
Pers. Fin. Controls	No	No	No	No	No	Yes	Yes
Heard News Controls	No	No	No	No	No	No	Yes

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is the 6-month change in the household's 1-year ahead inflation expectations. All independent variables are in first-differences (i.e., 6-month changes over the household's survey response) as in equation (12). "Demographic controls" include gender, age, education level, marital status, having children, and income. "Personal finance controls" include dummies representing whether the respondent's personal finances are better or worse than a year ago, and whether they expect their personal finances to be better or worse a year from now. "Heard News Controls" include dummies indicating whether the respondent reported hearing of plants opening or closing, hearing the stock market is rising or declining, hearing improvements or declines in specific industries, hearing that credit is easy to get, hearing that interest rates are high, and hearing that employment is high or low. Two-way clustered standard errors (by household and time) are in parenthesis. * p<0.1, ** p <0.05, *** p < 0.01

Table 5: Treatment Differentiated by Response Order

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Heard unfavorable news (1st response)	0.543***	0.387***	0.372***	0.224***	0.203**	0.195**	0.191**
	(0.0940)	(0.0877)	(0.0877)	(0.0838)	(0.0837)	(0.0834)	(0.0831)
Heard favorable news (1st response)	-0.534***	-0.441***	-0.439***	-0.343**	-0.354**	-0.341**	-0.331**
	(0.139)	(0.138)	(0.138)	(0.137)	(0.139)	(0.139)	(0.138)
Heard unfavorable news (2nd response)	0.614***	0.482***	0.470***	0.381***	0.391***	0.379***	0.367***
	(0.0990)	(0.0957)	(0.0962)	(0.0936)	(0.0962)	(0.0958)	(0.0956)
Heard favorable news (2nd response)	-0.215*	-0.133	-0.138	-0.0554	-0.0927	-0.100	-0.0917
	(0.128)	(0.126)	(0.127)	(0.127)	(0.128)	(0.128)	(0.128)
Headline inflation (y/y)		0.269***					
		(0.0280)					
Core Inflation (y/y)			0.110**				
			(0.0450)				
Food inflation (y/y)			0.116***				
			(0.0227)				
Energy inflation (y/y)			0.0244***				
			(0.00342)				
Observations	67833	67833	67833	67833	65665	65665	65665
Adjusted R^2	0.002	0.009	0.010	0.000	0.002	0.003	0.004
$\lvert$unfav(1st)$\rvert$ = $\lvert$fav(1st)$\rvert$, pval	0.958	0.745	0.687	0.461	0.351	0.367	0.383
$\lvert$unfav(2nd)$\rvert$ = $\lvert$fav(2nd)$\rvert$, pval	0.0117	0.0261	0.0359	0.0376	0.0579	0.0749	0.0799
Time Fixed Effects	No	No	No	Yes	Yes	Yes	Yes
Demographic Controls	No	No	No	No	Yes	Yes	Yes
Pers. Fin. Controls	No	No	No	No	No	Yes	Yes
Heard News Controls	No	No	No	No	No	No	Yes

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is the 6-month change in the household's 1-year ahead inflation expectations. All independent variables are in first-differences (i.e., 6-month changes over the household's survey response) as in equation (12). "Demographic controls" include gender, age, education level, marital status, having children, and income. "Personal finance controls" include dummies representing whether the respondent's personal finances are better or worse than a year ago, and whether they expect their personal finances to be better or worse a year from now. "Heard News Controls" include dummies indicating whether the respondent reported hearing of plants opening or closing, hearing the stock market is rising or declining, hearing improvements or declines in specific industries, hearing that credit is easy to get, hearing that interest rates are high, and hearing that employment is high or low. Two-way clustered standard errors (by household and time) are in parenthesis. * p<0.1, ** p <0.05, *** p < 0.01

favorable news decreases inflation expectations by about 0.3 percentage points. As in Table 4, p-values testing the null hypothesis of symmetric news effects are shown at the bottom of the table. For all specifications, the high p-values indicate that one cannot reject the null hypothesis of symmetry for the first responses. Interpreted through the lens of the model, these patterns suggest that those who report inflation news as their first answer to the news question are more likely to have specifically sought out that inflation news, and therefore less likely to have been influenced by news selection.

The coefficients on the second responses are more positive than the coefficients on the first responses. A more positive coefficient on hearing favorable news implies a larger asymmetry because the first-response coefficient is negative, exactly as suggested by our structural model. Column (1) shows that a household that reports hearing unfavorable inflation news as their second response increases inflation expectations by about 0.6 percentage points. A report of hearing favorable news decreases inflation expectations by about 0.2 percentage points. Columns (2) to (7) show that adding controls attenuates the size of the coefficients. In the model with full controls, column (7), the coefficient on unfavorable inflation news shrinks to about 0.4 but remains highly statistically significant, while the coefficient on favorable news shrinks to less than 0.1 and is statistically insignificant. The robustly low p-values on the symmetry tests provide evidence of asymmetric effects of favorable and unfavorable news that corroborate the implications of the theoretical model.[16]

In addition to corroborating the implications of the theoretical model, these results offer indirect support for our presumption that, conditional on time and individual fixed-effects, inflation-seeking behavior is as-good-as-random. A natural alternative hypothesis is that households search for inflation news when they already hold pessimistic priors about inflation. We could capture this in our model by assuming a positive correlation between inflation curiosity shocks and an individual's priors about inflation.[17] If the data were generated from such a model, the first-response coefficients for both types of inflation news would shift upwards. Specifically, we would expect a larger absolute coefficient on unfavorable news and a smaller absolute coefficient on favorable news for first responses as well as second responses. However, our estimated coefficients for first responses are statistically indistinguishable, which is more consistent with the view that these individuals do not systematically hold prior beliefs of a particular type at the time of exposure to inflation news.

[16] Including the household's 1-year ahead expectation of gas price growth has a negligible effect on the estimates, as shown in Appendix Table A1.This survey question is not available for the full sample period—it began in May 2006. Columns (1) and (2) of of table A1 show that the sample decreases from 65,650 observations to 38,730.

[17] Recall that we included time and individual fixed effects in our empirical specification, so only a time-varying and individual-specific correlation would influence this correlation.

5.2 Comparison with Model

We now ask whether our simple model, which has been calibrated only to match the unconditional properties in the aggregate data on hearing inflation news, gives qualitatively reasonable effects regarding the exposure to news. To do this, we perform the regressions specified in equations (11) and (13) on model simulated data.[18] Since the inference problem solved by households in our model is non-standard, we do not have a closed form for agent's expectations conditional on their information. We therefore solve the model numerically by simulating a long sequence of inflation and a large cross section of households and estimating a flexible functional form that relates realized inflation with the information set of the household.[19]

When we estimate (11) on simulated data, we find $\hat{\beta}^{a,model} = 0.87$ and $\hat{\gamma}^{a,model} = -0.78$ for first responses and $\hat{\beta}^{b,model} = 1.84$ and $\hat{\gamma}^{b,model} = 0.11$ for second responses. These estimates are qualitatively consistent with our OLS results in Column (7) of Table 5, implying that (unintentional) exposure to unfavorable inflation drives an even larger increase inflation expectations, while unintentional exposure to positive inflation news has essentially no effect on inflation expectations. When instead we estimate the pooled specification, equation (13), we find $\hat{\beta}^{pooled,model} = 1.39$ and $\hat{\gamma}^{pooled,model} = -0.46$. This pattern for the pooled regression coefficients, essentially a weighted average of the first and second response coefficients, is also qualitatively quite consistent with the empirical patterns.

Of course, our model is stylized in many dimensions, and the size of the coefficients generated by the model is somewhat larger than in the data. One likely source of this difference is the fact that real-world inflation is persistent, meaning that households are likely to rely on prior information in computing their inflation expectations as well as current news. Since our model assumes i.i.d. inflation, allowing for this additional reliance on prior information would reduce the overall magnitude of expectations to news media.

5.3 A Back-of-the Envelope Calculation Over the Pandemic Period

Household inflation expectations reached a peak of 5.4% in March 2022, having increased 3.2 pp from the onset of the pandemic in March 2020. Over the same period, the share of

[18] Since the model-generated data are also i.i.d., it also makes no difference whether we run the regression in levels or first-differences.

[19] For each agent type, $i \in \mathcal{A}$ or $i \in \mathcal{B}$, we regress the inflation realization on a fifth-order polynomial in the household's signal. We conclude this degree of richness is sufficient because allowing for additional terms in the polynomials has no effects on household expectations. For households of type $i \notin \mathcal{A} \cup \mathcal{B}$, we need only compute a conditional mean.

households hearing unfavorable inflation news increased by 40 percentage points. Although a comprehensive assessment of the role of media coverage in explaining the increase in expectations would require a structural analysis, we can use our estimated coefficients to conduct a back-of-the-envelope calculation, providing a rough estimate of its potential impact based on our findings.

Multiplying the causal effect from our preferred specification with the increased share of household exposed to news would imply that, at its peak, news coverage increased inflation expectations by $0.37 \times 0.4 \approx 0.15$ percentage points per month. This corresponds to roughly 4.5% of the cumulative increase in inflation expectations during the period. While it is not obvious how to accumulate these effects over time, one natural approach is to assume that they decay at the same rate as individual inflation expectations do on average in our sample.[20] Accumulating in this way, COVID-era media coverage of inflation could explain an increase of 0.6 of a percentage point in inflation expectations, or 18% of the total increase during that period. Hence, this exercise suggests that news media could account for somewhere between 4 and 18% of the increase in aggregate inflation expectations over this period.

6 Conclusion

This study points to an important role of news media in shaping households' inflation expectations. The model we develop shows that news selection leads to an asymmetry: high-inflation news will have a larger impact on expectations than low-inflation news. We estimate the effect of news on expectations using micro panel data, which allows for the inclusion of time and household fixed effects and to differentiate between different types of news exposure. We find that hearing news related to higher prices causes expectations to increase by around 0.4 percentage point. In line with the model, hearing news of lower prices has a near-zero effect. A back-of-the-envelope calculation shows that the media can play a substantial role in explaining inflation expectations.

[20]We estimate a AR(1) model coefficient of expectations in the panel of 0.40 over six months, which corresponds to a monthly autocorrelation coefficient of $0.40^{\frac{1}{6}} = 0.86$.

References

Afrouzi, H., A. Dietrich, K. Myrseth, R. Priftis, and R. Schoenle (2024, April). Inflation preferences. Working Paper 32379, National Bureau of Economic Research.

Armantier, O., S. Nelson, G. Topa, W. van der Klaauw, and B. Zafar (2016, 07). The price is right: Updating inflation expectations in a randomized price information experiment. *The Review of Economics and Statistics 98*(3), 503–523.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics 118*(1), 269–298.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017, July). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics 9*(3), 1–35.

Chahrour, R., K. Nimark, and S. Pitschner (2021). Sectoral media focus and aggregate fluctuations. *American Economic Review 111*(12), 3872–3922.

Coibion, O. and Y. Gorodnichenko (2015). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics 7*(1), 197–232.

Coibion, O., Y. Gorodnichenko, and M. Weber (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy 130*(6), 1537–1584.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021). Exposure to grocery prices and inflation expectations. *Journal of Political Economy 129*(5), 1615–1639.

Ehrmann, M., D. Pfajfar, and E. Santoro (2017). Consumer's attitudes and their inflation expectation. *International Journal of Central Banking 13*, 225.

Kmetz, A., A. H. Shapiro, and D. J. Wilson (2022). Can the news drive inflation expectations? *FRBSF Economic letter 2022*(31), 1–6.

Kuchler, T. and B. Zafar (2019). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance 74*(5), 2491–2542.

Larsen, V. H., L. A. Thorsrud, and J. Zhulanova (2021). News-driven inflation expectations and information rigidities. *Journal of Monetary Economics 117*, 507–520.

Malmendier, U. and S. Nagel (2016). Learning from inflation experiences. *The Quarterly Journal of Economics 131*(1), 53–87.

Nimark, K. P. (2014). Man-Bites-Dog Business Cycles. *American Economic Review 104*(8), 2320–67.

Nimark, K. P. and S. Pitschner (2019). News media and delegated information choice. *Journal of Economic Theory 181*, 160 – 196.

Shapiro, A. H. and D. J. Wilson (2022). Taking the fed at its word: A new approach to estimating central bank objectives using text analysis. *The Review of Economic Studies 89*(5), 2768–2805.

Trehan, B. et al. (2011). Household inflation expectations and the price of oil: it's déja vu all over again. *FRBSF Economic Letter 16*.

Weber, M., B. Candia, H. Afrouzi, T. Ropele, R. Lluberas, S. Frache, B. Meyer, S. Kumar, Y. Gorodnichenko, D. Georgarakos, O. Coibion, G. Kenny, and J. Ponce (2025). Tell me something i don't already know: Learning in low- and high-inflation settings. *Econometrica 93*(1), 229–264.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press.

Zaller, J. and S. Feldman (1992). A simple theory of the survey response: Answering questions versus revealing preferences. *American Journal of Political Science 36*(3), 579–616.

Online Appendix

A Proofs

To conserve notation it is useful to define the sets $\overline{\mathcal{A}_t}$ as set of agents who read a first story that is about above average inflation. More formally, $\overline{\mathcal{A}_t}$ contains the agent's index i if and only if $m_a(i) \in \Pi_t$ and $s_{m_a(i)} > \bar{\pi}$. Similarly, we can define $\underline{\mathcal{A}_t}$ as the set of agents who observe a first story about inflation that is (equal to or) below average. Evidently, $\mathcal{A}_t = \overline{\mathcal{A}_t} \cup \underline{\mathcal{A}_t}$. We later make use of analogous definitions of the sets $\overline{\mathcal{B}}_t$ and $\underline{\mathcal{B}}_t$.

Derivation of Corollary 1. Define the coefficients as in the main text, $\alpha^a \equiv \bar{\pi} = E[\pi^e_{it}|i \notin \mathcal{A}_t]$ and the slope coefficients are given by $\beta^a \equiv E[\pi^e_{it} - \bar{\pi}|i \in \overline{\mathcal{A}_t}]$, and $\gamma^a \equiv E[\pi^e_{it} - \bar{\pi}|\underline{\mathcal{A}_t}]$. There are three possible combinations of the dummy variables: whenever $i \in \overline{\mathcal{A}_t}$, then $U^a_{it} = 1$ and $F^a_{it} = 0$; whenever $i \in \underline{\mathcal{A}_t}$, then $U^a_{it} = 0$ and $F^a_{it} = 1$; and when whenever $i \notin \mathcal{A}_t$ then $U^a_{it} = F^a_{it} = 0$. We take these cases one at a time.

Case: $U^a_{it} = F^a_{it} = 0$: In this case, equation (5) becomes $\pi^e_{it} = \alpha^a + \mu^a_{it}$, which implies that $\mu^a_{it} = \pi^e_{it} - \alpha^a$. Use the definition of α^a and take conditional expectations of both sides of this to find:

$$\begin{aligned} E[\mu^a_{it}|i \notin \mathcal{A}_t] &= E[\pi^e_{it}|i \notin \mathcal{A}_t] - \alpha^a \\ &= E[\pi^e_{it}|i \notin \mathcal{A}_t] - E[\pi^e_{it}|i \notin \mathcal{A}_t] \\ &= 0. \end{aligned} \tag{16}$$

Case: $U^a_{it} = 1, F^a_{it} = 0$: In this case, equation (5) becomes $\pi^e_{it} = \alpha^a + \beta^a + \mu^a_{it}$, which implies that $\mu^a_{it} = \pi^e_{it} - \alpha^a - \beta^a$. Use the definition of α^a and β^a and take conditional expectations of both sides of this to find:

$$\begin{aligned} E[\mu^a_{it}|i \in \overline{\mathcal{A}_t}] &= E[\pi^e_{it}|i \in \overline{\mathcal{A}_t}] - \bar{\pi} - E[\pi^e_{it} - \bar{\pi}|i \in \overline{\mathcal{A}_t}] \\ &= E[\pi^e_{it} - \bar{\pi}|i \in \overline{\mathcal{A}_t}] - E[\pi^e_{it} - \bar{\pi}|i \in \overline{\mathcal{A}_t}] \\ &= 0. \end{aligned} \tag{17}$$

Case: $U^a_{it} = 0, F^a_{it} = 1$: In this case, equation (5) becomes $\pi^e_{it} = \alpha^a + \gamma^a + \mu^a_{it}$, which implies that $\mu^a_{it} = \pi^e_{it} - \alpha^a - \gamma^a$. Use the definition of α^a and γ^a and take conditional expectations of both sides of this to find:

$$
\begin{aligned}
E[\mu_{it}^a | i \in \underline{\mathcal{A}_t}] &= E[\pi_{it}^e | i \in \underline{\mathcal{A}_t}] - \bar{\pi} - E[\pi_{it}^e - \bar{\pi} | i \in \underline{\mathcal{A}_t}] \\
&= E[\pi_{it}^e - \bar{\pi} | i \in \underline{\mathcal{A}_t}] - E[\pi_{it}^e - \bar{\pi} | i \in \underline{\mathcal{A}_t}] \\
&= 0.
\end{aligned}
\tag{18}
$$

From (16) - (18), we immediately have that $E[\mu_{it}^a] = 0$. What remains to be shown is that $cov(\mu_{it}^a, U_{it}^a) = cov(\mu_{it}^a, F_{it}^a) = 0$. Compute $cov(\mu_{it}^a, U_{it}^a) = E[\mu_{it}^a U_{it}^a] - E[\mu_{it}^a]E[U_{it}^a]$. The second term is zero since, $E[\mu_{it}^a] = 0$. Expanding the first term, we have

$$
\begin{aligned}
E[\mu_{it}^a U_{it}^a] &= E[\mu_{it}^a U_{it}^a | i \in \overline{\mathcal{A}_t}] \times Pr[i \in \overline{\mathcal{A}_t}] + E[\mu_{it}^a U_{it}^a | i \notin \overline{\mathcal{A}_t}] \times Pr[i \notin \overline{\mathcal{A}_t}] \\
&= E[\mu_{it}^a | i \in \overline{\mathcal{A}_t}] \times Pr[i \in \overline{\mathcal{A}_t}] \times 1 + E[\mu_{it}^a | i \notin \overline{\mathcal{A}_t}] \times Pr[i \notin \overline{\mathcal{A}_t}] \times 0
\end{aligned}
$$

Both terms above zero: the first, since $E[\mu_{it}^a | i \in \overline{\mathcal{A}_t}] = 0$; the second, because the indicator function $U_{it}^a = 0$. Making the appropriate substitutions, the same steps establish that $cov(\mu_{it}^a, F_{it}^a) = 0$. $\square$

Proof of Proposition 1. For the proof of Propositions 1 - 3, we drop all time subscripts. Using Bayes law, compute the conditional distributions perceived by an agents i:

$$
f(\pi | i \in \overline{\mathcal{A}}) = \frac{Pr(s_{m_b(i),t} > \bar{\pi} | \pi, i \in \mathcal{A}) Pr(i \in \mathcal{A} | \pi) f(\pi)}{Pr(s_{m_b(i),t} > \bar{\pi} | i \in \mathcal{A}) Pr(i \in \mathcal{A})}
$$

$$
f(\pi | i \in \underline{\mathcal{A}}) = \frac{Pr(s_{m_b(i),t} \leq \bar{\pi} | \pi, i \in \mathcal{A}) Pr(i \in \mathcal{A} | \pi) f(\pi)}{Pr(s_{m_b(i),t} \leq \bar{\pi} | i \in \mathcal{A}) Pr(i \in \mathcal{A})}.
$$

Using the above expressions, we can compute

$$
E[\pi | i \in \overline{\mathcal{A}}] + E[\pi | i \in \underline{\mathcal{A}}] = \int \pi g(\pi) h(\pi) f(\pi) d\pi \tag{19}
$$

where $g(\pi) \equiv \left(\frac{ncdf(\pi - \bar{\pi})}{Pr(s_{m_b(i)} > \bar{\pi} | i \in \mathcal{A})} + \frac{1 - ncdf(\pi - \bar{\pi})}{Pr(s_{m_b(i)} \leq \bar{\pi} | i \in \mathcal{A})} \right)$, $ncdf(\cdot)$ is the normal cumulative distribution function for variance σ_η^2, and $h(\pi) \equiv \frac{Pr(i \in \mathcal{A} | \pi)}{Pr(i \in \mathcal{A})} = 1$. The symmetry of $f(\pi)$ and the signal noise imply that $Pr(s_{m_b(i)} \leq \bar{\pi} | i \in \mathcal{A}) = Pr(s_{m_b(i)} > \bar{\pi} | i \in \mathcal{A})$, so that the function $g(\pi)$ is also constant. We therefore have $E[\pi | i \in \overline{\mathcal{A}}] + E[\pi | i \in \underline{\mathcal{A}}] = 2\bar{\pi}$, which implies that

$$
E[\pi | i \in \overline{\mathcal{A}}] - \bar{\pi} = -\left(E[\pi | i \in \underline{\mathcal{A}}] - \bar{\pi} \right). \tag{20}
$$

Using the law of iterated expectations, we can replace π above with π_i^e, delivering (4) and completing the proof.

$\square$

In order to prove Proposition 2, we first need to establish the following Lemma.

Lemma 1. *Assume that*

1. $f(x)$ *is a distribution;*
2. $\tilde{f}(x) \equiv g(x)f(x)$ *is also distribution; and*
3. $g(x)$ *is a strictly increasing function, with* $E_f[g(x)] = 1$.

Then, $E_{\tilde{f}}[x] > E_f[x]$.

Proof of Lemma 1. Compute

$$\begin{aligned} E_{\tilde{f}}[x] - E_f[x] &= \int xg(x)f(x)dx - \int xf(x)dx \\ &= \int x(g(x)-1)f(x)dx \\ &= E_f[x(g(x)-1)] \\ &> E_f[g(x)-1]E_f[x] \\ &= 0. \end{aligned}$$

The inequality above comes from Jensen's inequality, since both $h(x) = x$ and $m(x) = g(x) - 1$ are strictly increasing in x. $\square$

Proof of Proposition 2. The first inequality in equation (6) follows from applying Bayes theorem,

$$f(\pi|i \in \mathcal{B}) = \frac{Pr(i \in \mathcal{B}|\pi)f(\pi)}{Pr(i \in \mathcal{B})}. \tag{21}$$

The ratio $\frac{Pr(i\in\mathcal{B}|\pi)}{Pr(i\in\mathcal{B})} = \frac{\Phi(\pi)}{Pr(i\in\mathcal{B})}$ is strictly increasing since $\Phi(\cdot)$ is strictly increasing. Applying Lemma 1 and using the law of iterated expectations to replace π with π_i^e gives the result. An analogous argument establishes the second inequality. $\square$

Proof of Proposition 3. Using Bayes law, compute the conditional distributions perceived by an agents i:

$$f(\pi|i \in \overline{\mathcal{B}}) = \frac{Pr(s_{m_b(i),t} > \bar{\pi}|\pi, i \in \mathcal{B})Pr(i \in \mathcal{B}|\pi)f(\pi)}{Pr(s_{m_b(i),t} > \bar{\pi}|i \in \mathcal{B})Pr(i \in \mathcal{B})}$$

$$f(\pi|i \in \underline{\mathcal{B}}) = \frac{Pr(s_{m_b(i),t} \le \bar{\pi}|\pi, i \in \mathcal{B})Pr(i \in \mathcal{B}|\pi)f(\pi)}{Pr(s_{m_b(i),t} \le \bar{\pi}|i \in \mathcal{B})Pr(i \in \mathcal{B})}.$$

Using the above expressions, we can compute

$$E[\pi|i \in \overline{\mathcal{B}}] + E[\pi|i \in \underline{\mathcal{B}}] = \int \pi g(\pi) h(\pi) f(\pi) d\pi \tag{22}$$

where $g(\pi) \equiv \left(\frac{ncdf(\pi - \bar{\pi})}{Pr(s_{m_b(i)} > \bar{\pi}|i \in \mathcal{B})} + \frac{1 - ncdf(\pi - \bar{\pi})}{Pr(s_{m_b(i)} \leq \bar{\pi}|i \in \mathcal{B})} \right)$, $ncdf(\cdot)$ is the normal cumulative distribution function for variance σ_η^2, and $h(\pi) \equiv \frac{Pr(i \in \mathcal{B}|\pi)}{Pr(i \in \mathcal{B})}$. The symmetry of $f(\pi)$ and the signal noise imply that $Pr(s_{m_b(i)} \leq \bar{\pi}|i \in \mathcal{B}) = Pr(s_{m_b(i)} > \bar{\pi}|i \in \mathcal{B})$, so that the function $g(\pi)$ is constant. Meanwhile, $h(\pi)$ is increasing by the assumption about the news selection function. Using Jensen's inequality, we therefore have

$$E[\pi|i \in \overline{\mathcal{B}}] + E[\pi|i \in \underline{\mathcal{B}}] > \bar{\pi} E[g(\pi)] E[h(\pi)] = 2\bar{\pi} \tag{23}$$

since $E[g(\pi)] = 2$ and $E[h(\pi)] = 1$ by their definitions. However, by Proposition 2, $\bar{\pi} > E[\pi|i \notin \mathcal{B}]$, so that (23) implies $E[\pi|i \in \overline{\mathcal{B}}] + E[\pi|i \in \underline{\mathcal{B}}] > 2E[\pi|i \notin \mathcal{B}]$. Rearranging, we have $E[\pi|i \in \overline{\mathcal{B}}] - E[\pi|i \notin \mathcal{B}] > -\left(E[\pi|i \in \underline{\mathcal{B}}] - E[\pi|i \notin \mathcal{B}]\right)$. Since $E[\pi|i \in \overline{\mathcal{B}}] > E[\pi|i \in \underline{\mathcal{B}}]$, we can conclude

$$|E[\pi|i \in \overline{\mathcal{B}}] - E[\pi|i \notin \mathcal{B}]| > |E[\pi|i \in \underline{\mathcal{B}}] - E[\pi|i \notin \mathcal{B}]|. \tag{24}$$

Using the law of iterated expectations, we can replace π in (24) with π_i^e, delivering (7) and completing the proof.

□

Derivation of Corollaries 2-3. Corollary 2 can be derived using steps virtually identical to the steps used to derive Corollary 1. Meanwhile, Corollary 3 follows by combining the representations of expectations for each type of agent described by Corollaries 1 and 2. □

Derivation of equations (14) - (15). The econometrician estimates

$$\pi_{it}^e = \alpha + \beta U_{it} + \gamma F_{it} + \tilde{\mu}_{it}. \tag{25}$$

We compute the expected cross product of the regressors:

$$E([1, U_{it}, F_{it}][1, U_{it}, F_{it}]') = \begin{pmatrix} 1 & \mu^U & \mu^F \\ \mu^U & \mu^U & 0 \\ \mu^F & 0 & \mu^F \end{pmatrix} \tag{26}$$

Here, we have defined $\mu^U \equiv E[U_{it}]$ and $\mu^F \equiv E[F_{it}]$ and observed that $E[U_{it}^2] = E[U_{it}]$,

$E[F_{it}^2] = E[F_{it}]$, and $E[U_{it}F_{it}] = 0$.

Using the data generating process in (10), we can compute

$$E[[1, U_{it}, F_{it}] \times \pi_t]' = \begin{pmatrix} \alpha^b + \beta^b\mu^U + \gamma^b\mu^f + \frac{1}{2}\phi(\beta^a - \beta^b + \Delta\alpha) + \frac{1}{2}\phi(\gamma^a - \gamma^b + \Delta\alpha) \\ \alpha^b\mu^U + \beta^b\mu^U + \frac{1}{2}\phi(\beta^a - \beta^b + \Delta\alpha) \\ \alpha^b\mu^F + \gamma^b\mu^F + \frac{1}{2}\phi(\gamma^a - \gamma^b + \Delta\alpha) \end{pmatrix}, \tag{27}$$

where we have used the fact that

$$\begin{aligned} E[c_{it}U_{it}] &= E[c_{it}U_{it}^a] \\ &= E[c_{it}] \times E[\mathbf{1}(\pi_t + \eta_{mt} > \bar{\pi})] \\ &= \phi \times \frac{1}{2}. \end{aligned}$$

Using these results, we compute the (population) regression coefficients

$$\begin{aligned} [\hat{\alpha}, \hat{\beta}, \hat{\gamma}] &= E[[1, U_{it}, F_{it}] \times \pi]E([1, U_{it}, F_{it}][1, U_{it}, F_{it}]')^{-1} \\ &= \left[\alpha^b, \beta^b + \frac{\phi}{2}\frac{\beta^a - \beta^b + d\alpha}{\mu^U}, \gamma^b + \frac{\phi}{2}\frac{\gamma^a - \gamma^b + d\alpha}{\mu^U}\right]. \end{aligned}$$

Finally, observe that by the definition of the coefficients, $\beta^a - \beta^b + d\alpha = E[\pi_{it}^e | i \in \overline{\mathcal{A}_t}] - E[\pi_{it}^e | i \in \overline{\mathcal{B}_t}]$, which is clearly negative. Similarly, $\gamma^a - \gamma^b + d\alpha = E[\pi_{it}^e | i \in \underline{\mathcal{A}_t}] - E[\pi_{it}^e | i \in \underline{\mathcal{B}_t}]$, which is also negative.

□

B Additional Materials

Figure A1: Inflation expectations versus newspaper coverage of inflation

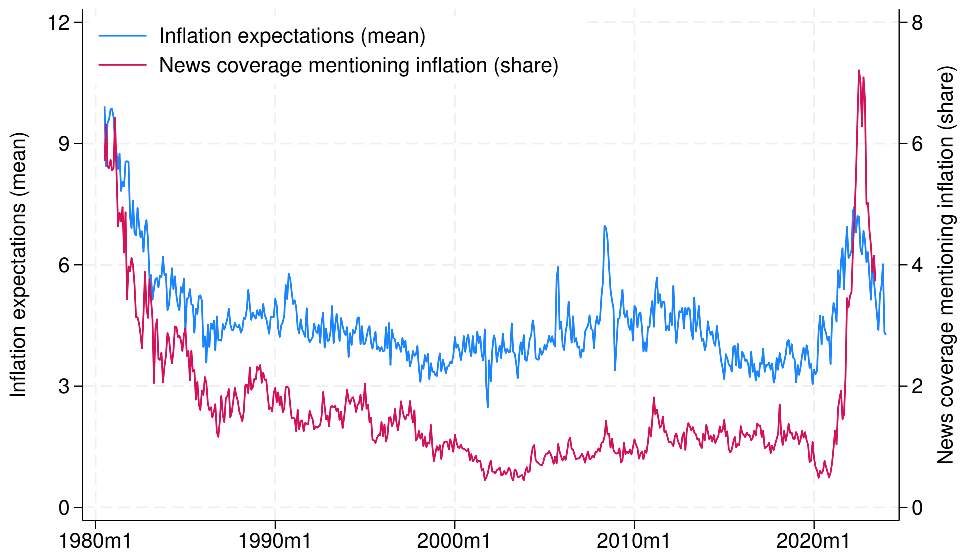

Figure A2: University of Michigan survey question regarding "heard news"

A6. During the last few months, have you heard of any favorable or unfavorable changes in business conditions?

YES	NO; HAVEN'T HEARD
	GO TO A7

A6a. What did you hear? (Have you heard of any other favorable or unfavorable changes in business conditions?)

IF NOT CLEAR WHETHER A CHANGE IS FAVORABLE OR UNFAVORABLE, PROBE: "Would (MENTION CHANGE) be favorable or unfavorable?" AND NOTE "FAVORABLE" OR "UNFAVORABLE."

Figure A3: University of Michigan survey codebook

NEWS	"During the last few months, have you heard of any favorable or unfavorable changes in business conditions?"	
	NEWS_F	Favorable news
	NEWS_U	Unfavorable news
	NEWS_N	No news
	NEWS_R	FAVORABLE - UNFAVORABLE + 100
NEWSRN	"What did you hear?"	
	NEWSRN_F_GOVT	Favorable news: Government; elections
	NEWSRN_F_EMP	Favorable news: Employment
	NEWSRN_F_DEM	Favorable news: High consumer demand
	NEWSRN_F_PRI	Favorable news: Lower prices
	NEWSRN_F_CRED	Favorable news: Easier credit
	NEWSRN_F_STK	Favorable news: Stock Market
	NEWSRN_F_TRD	Favorable news: Trade deficit
	NEWSRN_U_GOVT	Unfavorable news: Government; Elections
	NEWSRN_U_EMP	Unfavorable news: Unemployment
	NEWSRN_U_DEM	Unfavorable news: Lower consumer demand
	NEWSRN_U_PRI	Unfavorable news: Higher prices
	NEWSRN_U_CRED	Unfavorable news: Tighter credit
	NEWSRN_U_ENG	Unfavorable news: Energy crisis
	NEWSRN_U_STK	Unfavorable news: Stock Market
	NEWSRN_U_TRD	Unfavorable news: Trade deficit
	NEWSRN_NP	NEWSRN_F_PRI - NEWSRN_U_PRI
	NEWSRN_NE	NEWSRN_F_EMP - NEWSRN_U_EMP
	NEWSRN_NG	NEWSRN_F_GOVT - NEWSRN_U_GOVT

Table A1: Panel Estimates, including gas price survey control

	(1)	(2)	(3)	(4)
Heard unfavorable news (1st response)	0.203**	0.169*	0.146	0.133
	(0.0837)	(0.0996)	(0.101)	(0.100)
Heard favorable news (1st response)	-0.354**	-0.353**	-0.351**	-0.329*
	(0.139)	(0.170)	(0.169)	(0.168)
Heard unfavorable news (2nd response)	0.391***	0.349***	0.334***	0.310***
	(0.0962)	(0.112)	(0.110)	(0.110)
Heard favorable news (2nd response)	-0.0927	0.0400	0.0386	0.0290
	(0.128)	(0.169)	(0.170)	(0.170)
% Gas Prices up next yr			0.00849***	0.00828***
			(0.000922)	(0.000916)
Observations	65665	38744	38744	38744
Adjusted R^2	0.002	0.002	0.006	0.009
$\lvert$unfav(1st)$\rvert$ = $\lvert$fav(1st)$\rvert$, pval	0.351	0.356	0.304	0.322
$\lvert$unfav(2nd)$\rvert$ = $\lvert$fav(2nd)$\rvert$, pval	0.0579	0.0519	0.0633	0.0922
Time Fixed Effects	Yes	Yes	Yes	Yes
Demographic Controls	Yes	Yes	Yes	Yes
Pers. Fin. Controls	No	No	No	Yes
Heard News Controls	No	No	No	Yes

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is the 6-month change in the household's 1-year ahead inflation expectations. All independent variables are in first-differences (i.e., 6-month changes over the household's survey response) as in equation (12). Two-way clustered standard errors (by household and time) are in parenthesis. * p<0.1, ** p <0.05, *** p < 0.01