NBER WORKING PAPER SERIES

THE PERCEIVED INFLATION WEDGE

Neville Francis

Working Paper 35354
http://www.nber.org/papers/w35354

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2026

I thank Valerie Ramey and Wayne Robinson for insightful discussions. I am indebted to Rolando Rodriguez and Matthew Jansen for their assistance with deriving the media mention and search activity portions of the perceived inflation index. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The Perceived Inflation Wedge
Neville Francis
NBER Working Paper No. 35354
June 2026
JEL No. C32, D84, E21, E31, E71

ABSTRACT

Standard measures of inflation, such as the Consumer Price Index (CPI), are designed to measure changes in the cost of living, not how households perceive price changes under limited attention and imperfect information. This paper introduces Perceived Inflation (PI), an attention-weighted inflation measure derived from a model in which households allocate attention across goods and observe noisy price signals. I construct monthly PI using disaggregated CPI data and attention proxies based on price salience, volatility, media mentions, and search behavior. The resulting PI series is highly correlated with CPI inflation but exhibits greater short-run volatility, especially during the Global Financial Crisis, the COVID-19 recession, and the 2021–22 inflation surge. A state-space version of the index implies a persistent latent attention component and a time-varying PI wedge. Finally, the paper embeds the estimated PI wedge into a heterogeneous-household model to demonstrate how attention-driven distortions in inflation perceptions affect saving and consumption through perceived real returns.

Neville Francis
University of North Carolina at Chapel Hill
Department of Economics
and NBER
nfrancis@unc.edu

1 Introduction

Official measures of inflation, such as the Consumer Price Index (CPI), are designed to measure changes in the cost of living using fixed expenditure weights. Yet households do not experience inflation through a fully representative basket, nor do they process all price changes equally. A growing body of evidence shows that consumers disproportionately attend to prices that are salient, volatile, frequently encountered, or amplified through media and search activity. As a result, inflation as measured statistically can differ systematically from inflation as perceived by households.

This paper attempts to quantify households' attention to prices. The aim is to show how limited attention distorts inflation perceptions, how those distortions shape expectations, and how they affect household behavior. I argue that perception and expectation formation are distinct stages: attention distorts perceived inflation by altering which prices households notice, while expectation formation selectively filters these distortions depending on the informational environment.

To formalize this mechanism, I introduce a novel measure of inflation—*Perceived Inflation* (PI)—and develop a microfounded model in which households face noisy information about individual prices with signal-acquisition. In the model, households optimally allocate scarce attention across goods, generating endogenous and time-varying attention weights. Perceived inflation then emerges as an attention-weighted aggregation of price changes, producing a structural wedge between itself and the official measure of inflation.

Using disaggregated CPI category data and proxies for household attention, I construct a monthly measure of perceived inflation (PI). Although PI closely tracks CPI inflation, it exhibits substantially greater short-run volatility, particularly during periods of heightened macroeconomic stress. I then use the estimated perceived-inflation wedge to illustrate one potential macroeconomic application by embedding it in a perception-augmented heterogeneous-agent model. In this framework, households make intertemporal decisions based on perceived rather than official inflation, so deviations between perceived and realized inflation distort perceived real returns, saving behavior, and consumption choices. The findings demonstrate that attention-driven distortions in inflation perceptions can have meaningful effects on household balance sheets and consumption outcomes. Variation in the perceived-inflation wedge alters wealth accumulation, the fraction of hand-to-mouth households, and marginal propensities to consume, with the largest effects concentrated among low-wealth and liquidity-constrained households.

The remainder of the paper proceeds as follows. Section 2 provides empirical motivation and reviews related theoretical frameworks. Section 3 develops the model and constructs the perceived inflation index. Section 4 details the construction of the attention weights and perceived inflation. Section 5 explores one potential macroeconomic application of the perceived-inflation wedge by embedding it into a heterogeneous-household model and examining its implications for household behavior. Finally, Section 6 concludes.

Related Literature: This paper contributes to a growing literature showing that households' inflation perceptions systematically differ from official inflation measures. Early work by Jonung (1981) and Detmeister et al. (2016) documents that households tend to perceive inflation as higher than measured CPI inflation, particularly for frequently purchased goods. Experimental and household-level evidence further shows that consumers overweight salient and personally experienced price changes when forming inflation beliefs and expectations (Georganas et al. (2014); D'Acunto et al. (2021); Kuchler and Zafar (2019); Malmendier and Nagel (2016)).

The approach in this paper is also related to recent work emphasizing the role of attention and information frictions in expectation formation. D'Acunto et al. (2023) and Weber et al. (2022) highlight the importance of salience, disagreement, and subjective inflation experiences in shaping household expectations and macroeconomic dynamics. While this literature documents that inflation perceptions differ from official inflation, less is known about how those distortions are transmitted into expectations and household decisions. Most closely related is Pfäuti (2026), which shows that attention to inflation is state-dependent and rises sharply during high-inflation periods, amplifying the transmission of supply shocks.

2 Motivation

Households' perceptions of inflation often diverge from official measures such as the Consumer Price Index (CPI). While CPI is designed to track changes in the cost of living for a representative consumption basket, it is not intended to measure how households actually notice and internalize price changes in daily life. This distinction motivates the central question: why can inflation as noticed by households differ from inflation as measured statistically? I introduce Perceived Inflation (PI) as a framework for addressing this gap. CPI measures the cost of living; PI measures the cost of noticing prices.

2.1 Theoretical Foundations

The divergence between measured and perceived inflation can be understood through models of limited attention and bounded information processing. The PI reflects this divergence by allowing inflation weights to vary with attention rather than remaining fixed at expenditure shares. I focus on three complementary frameworks with related mechanisms: rational inattention, salience, and sparsity-based attention.

2.1.1 Rational Inattention (Sims (2003))

Rational inattention (RI) theory posits that individuals face cognitive constraints and optimally allocate attention subject to an information-processing cost (Sims (2003)). Consumers therefore do not process all price changes equally, but instead devote greater attention to prices that are more informative or economically important.

The PI can be motivated by a rational inattention problem of the form:

$$\max_{\{\omega_{i,t}\}} \mathbb{E}[U(\mathbf{p}_t)] - \eta \cdot I(\mathbf{p}_t; \mathbf{s}_t),$$

where $I(\cdot)$ denotes mutual information between the true price vector $\mathbf{p}_t$ and the observed signal $\mathbf{s}_t$, and η governs the marginal cost of attention. In this framework, the attention weights $\omega_{i,t}$ emerge endogenously from the tradeoff between information acquisition and cognitive cost.

2.1.2 Salience Theory (Bordalo et al. (2013))

Salience theory emphasizes that consumers overweight price changes that are unusually large or prominent relative to their surrounding context (Bordalo et al. (2013)). This implies that perceived inflation places disproportionate weight on goods with unusually large relative price movements. A salience-adjusted weighting scheme takes the form:

$$\omega_{i,t} \propto |\pi_{i,t} - \bar{\pi}_t|,$$

where $\bar{\pi}_t$ denotes average inflation across goods.

2.1.3 Sparsity-Based Attention (Gabaix (2014))

The sparse-max framework of Gabaix (2014) provides a complementary behavioral foundation. Rather than processing all available information, agents selectively attend to variables that are most important for decision-making.

Let $\pi_{i,t}$ denote the price change of good i, and let $\theta_{i,t} \in [0,1]$ represent the attention allocated to that good. Perceived price changes are given by

$$\pi^s_{i,t} = \theta_{i,t}\pi_{i,t} + (1-\theta_{i,t})\bar{\pi}_{i,t},$$

where $\bar{\pi}_{i,t}$ is a default or anchor value.

Attention is costly, and households choose attention intensities to minimize perception loss together with a sparsity penalty:

$$\min_{\theta\in[0,1]^n} \frac{1}{2}\sum_{i=1}^{n}(1-\theta_{i,t})^2\sigma^2_{\pi_i}\left(\frac{\partial U}{\partial x_{i,t}}\frac{\partial x_{i,t}}{\partial p_{i,t}}\right)^2 + \lambda\sum_{i=1}^{n}|\theta_{i,t}|.$$

This formulation implies that attention is concentrated on goods whose price changes are more consequential for utility. In the empirical perceived-inflation index developed below, these attention intensities motivate normalized aggregation weights of the form

$$\omega_{i,t} = \frac{w_i^{exp}\theta_{i,t}}{\sum_{j=1}^{N} w_j^{exp}\theta_{j,t}},$$

so that

$$\pi_t^{PI} = \sum_{i=1}^{N}\omega_{i,t}\pi_{i,t}.$$

3 A Model of Attention, Information, and Perceived Inflation

This section develops a simple model in which perceived inflation arises endogenously from optimal attention allocation. Time is discrete, $t = 0, 1, 2, \ldots$. A continuum of goods $i \in [0, 1]$ is available each period, each with nominal price $p_{i,t}$. Households are subject to limited attention and noisy information about individual price changes.

3.1 Preferences and Consumption Aggregation

Households consume a CES basket of goods,

$$C_t = \left(\int_0^1 c_{i,t}^{\frac{\sigma-1}{\sigma}} \, di \right)^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 0, \tag{1}$$

and maximize expected discounted utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(C_t), \qquad u'(C) > 0, \; u''(C) < 0, \tag{2}$$

subject to a standard budget constraint. The ideal CES price index is

$$P_t = \left(\int_0^1 p_{i,t}^{1-\sigma} \, di \right)^{\frac{1}{1-\sigma}}, \tag{3}$$

and the corresponding aggregate inflation rate is

$$\pi_t \equiv \log P_t - \log P_{t-1}. \tag{4}$$

Item-level inflation is defined by

$$\pi_{i,t} \equiv \log p_{i,t} - \log p_{i,t-1}. \tag{5}$$

Let $w_i^{\exp}$ denote the CPI expenditure weights, so that $\int_0^1 w_i^{\exp}\, di = 1$ and we can write true inflation as

$$\pi_t = \int_0^1 w_i^{\exp}\, \pi_{i,t}\, di. \tag{6}$$

3.2 Information and Signal-Acquisition Technology

Households do not observe $\pi_{i,t}$ directly. Instead, they observe noisy signals

$$s_{i,t} = \pi_{i,t} + \varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim \mathcal{N}\big(0,\, \sigma^2_{\varepsilon,i}(\ell_{i,t})\big)\,, \tag{7}$$

where $\ell_{i,t} \geq 0$ denotes information-processing effort devoted to good i at time t. This effort improves the precision of the signal about good i's price change. Higher information-processing effort lowers signal noise according to

$$\sigma^2_{\varepsilon,i}(\ell_{i,t}) = \frac{\bar{\sigma}_i^2}{1 + \phi \ell_{i,t}}, \qquad \phi > 0. \tag{8}$$

Attention operates through costly signal acquisition. The per-period information-processing cost is

$$\mathcal{I}_t = \int_0^1 \kappa(\ell_{i,t})\, di, \qquad \kappa(\ell) = \frac{\lambda}{2}\ell^2, \qquad \lambda > 0. \tag{9}$$

Lifetime utility is

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[u(C_t) - \mathcal{I}_t\right].$$

3.3 Bayesian Perception of Item-Level Inflation

Households hold Gaussian priors for item-level inflation, $\pi_{i,t} \sim \mathcal{N}(\bar{\pi}, \sigma^2_{\pi,i})$. Given $s_{i,t}$, the posterior mean is

$$\mathbb{E}_t[\pi_{i,t} \mid s_{i,t}] = \bar{\pi} + K_{i,t}\big(s_{i,t} - \bar{\pi}\big), \tag{10}$$

where the Kalman gain is

$$K_{i,t} = \frac{\sigma^2_{\pi,i}}{\sigma^2_{\pi,i} + \sigma^2_{\varepsilon,i}(\ell_{i,t})}, \qquad 0 < K_{i,t} < 1. \tag{11}$$

The corresponding posterior variance is

$$\mathrm{Var}(\pi_{i,t} \mid s_{i,t}) = (1 - K_{i,t})\sigma^2_{\pi,i}. \tag{12}$$

3.4 Perceived Inflation as an Attention-Weighted Aggregate

Given item-level signals, households form beliefs about individual price changes. Let $\mathbb{E}_t[\pi_{i,t} \mid s_{i,t}]$ denote the posterior expectation of inflation for good i based on the signal $s_{i,t}$. Perceived aggregate inflation is defined as the aggregation of these posterior means using expenditure weights:

$$\pi_t^{\mathrm{PI},\,post} \equiv \int_0^1 w_i^{\exp}\, \mathbb{E}_t[\pi_{i,t} \mid s_{i,t}]\, di. \tag{13}$$

Using the Kalman updating rule,

$$\mathbb{E}_t[\pi_{i,t} \mid s_{i,t}] = \bar{\pi} + K_{i,t}(s_{i,t} - \bar{\pi}), \tag{14}$$

and the signal equation $s_{i,t} = \pi_{i,t} + \varepsilon_{i,t}$, the posterior mean can be written as

$$\mathbb{E}_t[\pi_{i,t} \mid s_{i,t}] = \bar{\pi} + K_{i,t}(\pi_{i,t} - \bar{\pi}) + K_{i,t}\varepsilon_{i,t}. \tag{15}$$

Substituting into the aggregation yields

$$\pi_t^{\mathrm{PI},\,post} = \bar{\pi} + \int_0^1 w_i^{\exp} K_{i,t}(\pi_{i,t} - \bar{\pi})\, di + \int_0^1 w_i^{\exp} K_{i,t}\varepsilon_{i,t}\, di. \tag{16}$$

Taking expectations over the signal noise, the last term vanishes, and perceived inflation is centered around

$$\pi_t^{\mathrm{PI},\,post} = \bar{\pi} + \int_0^1 w_i^{\exp} K_{i,t}(\pi_{i,t} - \bar{\pi})\, di. \tag{17}$$

Rearranging terms gives

$$\pi_t^{\mathrm{PI},\,post} = \bar{\pi}\left(1 - \int_0^1 w_i^{\exp} K_{i,t}\, di\right) + \int_0^1 w_i^{\exp} K_{i,t}\pi_{i,t}\, di. \tag{18}$$

This representation shows that perceived inflation depends on expenditure weights and on signal precision through the Kalman gain. Since the Kalman gain is increasing in signal-acquisition effort $\ell_{i,t}$, households place greater effective weight on goods for which they acquire more precise price signals.

The effective terms $w_i^{\exp} K_{i,t}$ need not sum to one. Therefore, for empirical implementation, I construct a normalized perceived-inflation index that preserves the relative signal-precision distortions implied by the Kalman gains:

$$\pi_t^{\mathrm{PI}} \equiv \int_0^1 \omega_{i,t}\, \pi_{i,t}\, di, \tag{19}$$

where

$$\omega_{i,t} \equiv \frac{w_i^{\exp} K_{i,t}}{\int_0^1 w_j^{\exp} K_{j,t}\, dj}. \tag{20}$$

By construction,

$$\int_0^1 \omega_{i,t}\, di = 1. \tag{21}$$

The normalized weights $\omega_{i,t}$ therefore combine CPI expenditure shares with signal-extraction intensity. Goods with larger Kalman gains receive greater weight in perceived inflation relative to their CPI expenditure shares.

3.5 Perception Loss and the Role of Squared Weights

To make the role of expenditure weights clear, it is useful to spell out the loss from misperceiving aggregate inflation. For the purpose of deriving the signal-acquisition problem, I use the posterior perceived-inflation object defined above and define the posterior perception error as

$$e_t^{post} \equiv \pi_t^{PI,post} - \pi_t. \tag{22}$$

Using the definitions of $\pi_t^{PI,post}$ and π_t, we can write

$$\begin{aligned} e_t^{post} &= \int_0^1 w_i^{\exp}\, \mathbb{E}_t[\pi_{i,t} \mid s_{i,t}]\, di - \int_0^1 w_i^{\exp}\, \pi_{i,t}\, di & (23) \\ &= \int_0^1 w_i^{\exp} \big(\mathbb{E}_t[\pi_{i,t} \mid s_{i,t}] - \pi_{i,t}\big)\, di. & (24) \end{aligned}$$

The household cares about the expected squared posterior perception error,

$$\mathbb{E}_t\Big[(e_t^{post})^2\Big] = \mathbb{E}_t\left[\left(\int_0^1 w_i^{\exp} \big(\mathbb{E}_t[\pi_{i,t} \mid s_{i,t}] - \pi_{i,t}\big)\, di\right)^2\right]. \tag{25}$$

Assuming that the item-level perception errors are conditionally uncorrelated across goods,[1] the cross terms vanish and we obtain

$$\mathbb{E}_t\Big[(e_t^{post})^2\Big] = \int_0^1 (w_i^{\exp})^2 \,\mathrm{Var}\big(\mathbb{E}_t[\pi_{i,t} \mid s_{i,t}] - \pi_{i,t}\big)\, di. \tag{26}$$

Using the posterior variance from above,

$$\mathrm{Var}\big(\mathbb{E}_t[\pi_{i,t} \mid s_{i,t}] - \pi_{i,t}\big) = (1 - K_{i,t})\sigma_{\pi,i}^2, \tag{27}$$

we obtain

$$\mathbb{E}_t\Big[(e_t^{post})^2\Big] = \int_0^1 (w_i^{\exp})^2\, (1 - K_{i,t})\sigma_{\pi,i}^2\, di. \tag{28}$$

3.6 Optimal Signal-Acquisition Effort

I now embed the perception loss measure into the household's information-acquisition problem. Let the household's per-period loss from misperceiving inflation be proportional to the mean squared perception error,

$$L_t \equiv \frac{\gamma}{2}\,\mathbb{E}_t[e_t^{post}]^2], \qquad \gamma > 0. \tag{29}$$

[1]The expected squared posterior perception error contains both variance and covariance terms:

$$\mathbb{E}_t[(e_t^{post})^2] = \int_0^1 (w_i^{\exp})^2 \mathrm{Var}_{i,t}\, di + \int\limits_{\substack{0\le i\le 1\\ 0\le j\le 1\\ i\neq j}} w_i^{\exp} w_j^{\exp} \mathrm{Cov}_{ij,t}\, di\, dj.$$

For tractability, the baseline model assumes that item-level perception errors are conditionally uncorrelated across goods, so that the covariance terms vanish. This assumption isolates the role of item-specific signal acquisition and avoids introducing a covariance structure across goods. Allowing correlated perception errors would add covariance terms but would not change the central mechanism that greater signal-acquisition effort lowers item-level perception variance.

Using the expression above, this becomes

$$L_t = \frac{\gamma}{2} \int_0^1 (w_i^{\text{exp}})^2 \, (1 - K_{i,t}) \sigma_{\pi,i}^2 \, di. \tag{30}$$

The household chooses signal-acquisition effort $\{\ell_{i,t}\}$ to trade off this perception loss against information-processing costs. The static signal-acquisition problem at time t is

$$\min_{\{\ell_{i,t}\}} \left\{ \frac{\gamma}{2} \int_0^1 (w_i^{\text{exp}})^2 (1 - K_{i,t}) \sigma_{\pi,i}^2 \, di \; + \; \int_0^1 \kappa(\ell_{i,t}) \, di \right\}. \tag{31}$$

Because $K_{i,t}$ depends on $\ell_{i,t}$ through the signal variance $\sigma_{\varepsilon,i}^2(\ell_{i,t})$, the problem determines the optimal amount of signal-acquisition effort devoted to each good.

Focusing on a given good i, the first-order condition for optimal signal-acquisition effort $\ell_{i,t}$ is

$$\frac{\partial}{\partial \ell_{i,t}} \left[(w_i^{\text{exp}})^2 (1 - K_{i,t}) \sigma_{\pi,i}^2 \right] + \kappa'(\ell_{i,t}) = 0. \tag{32}$$

The first term captures the marginal benefit of signal acquisition for good i: more effort raises $K_{i,t}$, lowers $(1 - K_{i,t})$, and reduces perception loss. The second term, $\kappa'(\ell_{i,t}) = \lambda \ell_{i,t}$ under quadratic costs, is the marginal cost of effort. Optimal signal acquisition equates the marginal benefit of improved signal precision with its marginal cost.

Under the signal-acquisition technology

$$\sigma_{\varepsilon,i}^2(\ell) = \frac{\bar{\sigma}_i^2}{1 + \phi \ell}, \qquad \phi > 0,$$

and quadratic costs, we can characterize how optimal effort responds to volatility and expenditure shares.

Proposition 1. *Under the assumptions above, optimal signal-acquisition effort satisfies*

$$\frac{\partial \ell_{i,t}}{\partial \sigma^2_{\pi,i}} > 0, \qquad \frac{\partial \ell_{i,t}}{\partial w_i^{\exp}} > 0.$$

More volatile goods and goods with larger expenditure shares receive strictly greater signal-acquisition effort. See the appendix for the proof.

Because higher signal-acquisition effort raises the Kalman gain $K_{i,t}$, goods with greater volatility or larger expenditure shares also receive larger perceived-inflation weights $\omega_{i,t}$. This mechanism makes perceived inflation more sensitive to categories whose prices are both economically important and informationally valuable.

3.7 Information Environment, Media Exposure, and Search Intensity

The preceding section treats signal-acquisition effort as a household choice. In practice, the precision of price signals also depends on the external information environment. Media coverage and search activity can make some price changes easier to observe or process, even holding household effort fixed.

Let media exposure $\mathcal{M}_{i,t}$ and search intensity $\mathcal{S}_{i,t}$ improve signal precision:

$$\sigma^2_{\varepsilon,i}(\ell_{i,t}, \mathcal{M}_{i,t}, \mathcal{S}_{i,t}) = \frac{\bar{\sigma}_i^2}{1 + \phi \ell_{i,t} + \phi_M \mathcal{M}_{i,t} + \phi_S \mathcal{S}_{i,t}}. \tag{33}$$

Then

$$\frac{\partial K_{i,t}}{\partial \mathcal{M}_{i,t}} > 0, \qquad \frac{\partial K_{i,t}}{\partial \mathcal{S}_{i,t}} > 0. \tag{34}$$

Proposition 2. *Media exposure and search intensity raise signal precision and therefore* *increase the Kalman gain. Through this channel, goods receiving greater media coverage or search activity receive larger perceived-inflation weights $\omega_{i,t}$. See the appendix for a formal proof.*

4 Estimating Attention Weights and the PI

In the model, attention operates through signal acquisition. Households devote signal-acquisition effort $\ell_{i,t}$ to goods whose price changes they monitor more closely. This effort improves signal precision, raises the Kalman gain $K_{i,t}$, and ultimately increases the perceived-inflation weight $\omega_{i,t}$. Since $\ell_{i,t}$ is not directly observed, I will use observable proxies for attention and information exposure to estimate a latent attention state $A_{i,t}$. This latent state is then mapped into signal precision, the Kalman gain, and finally the perceived-inflation weights.

The empirical mapping can therefore be summarized as

$$\text{attention proxies} \longrightarrow A_{i,t} \longrightarrow \sigma^2_{\varepsilon,i,t} \longrightarrow K_{i,t} \longrightarrow \omega_{i,t} \longrightarrow \pi_t^{PI}.$$

4.1 Components of the Attention Index

We now discuss each of the observable components used to measure category-level attention. The data sources for the components are listed in Table 1. Each component captures a different dimension of attention. Salience and volatility are derived from category-level price changes and measure which prices are most likely to attract notice or require monitoring. Media mentions capture exposure to public information, while Google Trends search intensity captures active information seeking.

Driver	Description / Proxy	Data Sources
Salience	Magnitude of absolute or relative price changes, e.g., $\|\Delta\pi_{i,t}\|$, or deviation from aggregate inflation $\|\pi_{i,t} - \bar{\pi}_t\|$	BLS disaggregated CPI category data
Volatility	Rolling standard deviation of price changes over a specified window to capture price variability	BLS disaggregated CPI category data
Media Mentions	Frequency and intensity of news and social media coverage related to specific goods or categories	ProQuest TDM
Search Behavior	Volume of online search queries and information-seeking behavior	Google Trends

Table 1: Proxies for Latent Category-Level Attention

Salience: Absolute Price Change

The first measure of salience captures the degree of recent price movement for each good, independent of direction. It is defined as the absolute value of the monthly log price change (i.e., inflation rate) for good i at time t:

$$s_{i,t}^{\text{abs}} = |\pi_{i,t}| = |\log(P_{i,t}) - \log(P_{i,t-1})|$$

This measure highlights goods experiencing notable price changes — whether increases or decreases — under the assumption that large price swings are more likely to draw consumer attention.

Salience: Deviation from Average Inflation

An alternative salience measure captures how unusual a good's price change is relative to the broader inflation environment. For good i at time t, this salience is defined as:

$$s_{i,t}^{\text{dev}} = |\pi_{i,t} - \bar{\pi}_t|$$

where $\pi_{i,t} = \log(P_{i,t}) - \log(P_{i,t-1})$ is the monthly inflation rate, and $\bar{\pi}_t$ is the cross-good average inflation in month t.[2] This formulation emphasizes goods that stand out from the general inflation trend — either due to exceptional increases or decreases — and reflects the intuition that consumers are more likely to notice price changes that deviate from the norm.

Volatility

Volatility captures the perceived instability of a good's price over time. Following standard practice, I define it as the rolling standard deviation of log price changes for each good i. Specifically, for a given month t, I compute:

$$\text{Volatility}_{i,t} = \text{std}\left(\pi_{i,t-k+1}, \pi_{i,t-k+2}, \ldots, \pi_{i,t}\right)$$

where k is the size of the rolling window (e.g., $k = 12$ months).

While Bates and Gabor (1986) do not explicitly define or estimate volatility in the time-series sense, their findings highlight the behavioral consequences of price instability. Consumers demonstrated persistent misestimation of long-term inflation effects and showed difficulty reasoning about exponential price growth. These inconsistencies suggest that perceived volatility — i.e., chronic unpredictability in prices — may drive attentional bias even in low-inflation environments. This supports the idea that volatility, as distinct from

[2] To ensure robustness, $\bar{\pi}_t$ is only computed when a minimum number of goods (e.g., five) report non-missing price changes in month t. Otherwise, $s_{i,t}^{\text{dev}}$ is treated as missing.

momentary salience, contributes to ongoing consumer attention.

Distinguishing Salience from Volatility. Salience and volatility are applied to actual price changes. Although both measures shape consumer attention, they capture distinct mechanisms. *Salience* refers to the size or rarity of a *recent* price change that draws immediate notice, such as a sudden jump in grocery prices. It is inherently short-term and event-driven. In contrast, *volatility* captures the *persistent instability* of a good's price over time. A good with frequent but small price fluctuations (e.g., gasoline or airfare) may not have highly salient changes, but its unpredictability can sustain consumer attention. A good can be highly salient but stable (e.g., rent that rises sharply once a year), or highly volatile but rarely salient (e.g., niche electronics with noisy price series). Therefore, the empirical attention measurement system includes both dimensions as separate inputs, ensuring that the model reflects both instantaneous and longer-term sources of consumer attention.

Media Mentions

To capture price changes in public discourse, I include media mentions as a driver of attention. Following Carroll (2003), media coverage plays a central role in the diffusion of macroeconomic information and the formation of household beliefs. For each good i, I define a set of keyword phrases $\mathcal{K}_i$ that reflect how the good is discussed in news and media sources. These keywords are used to count the frequency of relevant media coverage over time. Table 2 provides examples of keywords used in the search for media mentions. Details of the method used are described in Appendix C.

Let $n_{i,t}^M$ denote the number of times any of the keywords in $\mathcal{K}_i$ is mentioned in media

Table 2: Expanded Media Search Terms for CPI Categories

CPI Category	**Suggested Search Terms for Media Mentions and Search Intensity ($\mathcal{K}_i$)**
Food	"food prices", "grocery prices", "cost of groceries"
Shelter	"rent prices", "housing market", "cost of living"
Fuels and utilities	"energy prices", "utility bills", "electricity cost"
Household furnishings and operations	"furniture prices", "appliance costs", "home goods prices"
Men's and boys' apparel	"men's clothing prices", "men's fashion cost", "cost of men's clothes", similar terms for boys
Women's and girls' apparel	"women's clothing prices", "women's fashion cost", "cost of women's clothes", similar terms for girls
Footwear	"shoe prices", "footwear costs", "cost of shoes"
Infants' and toddlers' apparel	"baby clothes prices", "infant clothing cost", "toddler apparel prices"
Jewelry and watches	"jewelry prices", "watch prices", "gold and silver prices"
Private transportation	"gasoline prices", "car prices", "vehicle costs"
Public transportation	"bus fare", "public transport cost", "subway prices"
Medical care commodities	"drug prices", "prescription costs", "medicine prices"
Medical care services	"doctor visit cost", "healthcare prices", "hospital bill"
Video and audio	"tv prices", "audio equipment prices", "electronics prices"
Pets, pet products and services	"pet food prices", "veterinary costs", "cost of owning a pet"
Sporting goods	"sporting goods prices", "fitness equipment prices", "gym gear cost"
Photography	"camera prices", "photography equipment cost", "digital camera prices"
Other recreational goods	"recreational equipment prices", "outdoor gear cost", "hobby goods prices"
Other recreation services	"recreation services cost", "amusement park prices", "sports tickets cost"
Recreational reading materials	"book prices", "magazine prices", "cost of reading materials"
Education	"college tuition", "education costs", "student loan burden"
Communication	"internet prices", "cell phone plan costs", "mobile data prices"
Tobacco and smoking products	"cigarette prices", "tobacco tax", "vape prices"
Personal care	"toiletry prices", "personal care products", "cost of grooming"

articles in month t. The media mentions component is then

$$\text{MediaMentions}_{i,t} = n^{M}_{i,t}.$$

These counts are typically obtained from structured media databases or APIs such as GDELT, LexisNexis, Media Cloud, ProQuest, or NewsAPI, using Boolean queries of the form:

```
"keyword1" OR "keyword2" OR ... "keywordN"
```

For example, the category *Private Transportation* might use the keyword set:

$$\mathcal{K}_i = \{\texttt{"gasoline prices"}, \texttt{"car prices"}, \texttt{"vehicle costs"}\}$$

To collect media mentions relevant to this study, I used ProQuest TDM Studio, a text and data mining solution that provides programmatic access to a wide array of licensed content across disciplines, including major newspapers, news, blogs, audio and visual, and historical primary sources. TDM Studio enables researchers to define and export curated corpora using keyword searches and filtering tools within its web-based interface. The resulting datasets were analyzed in a Jupyter environment with Python, which TDM Studio natively supports.

Data Alignment. Media mentions are aggregated to the monthly level and merged with the corresponding CPI category time series by calendar month.

Scaling by Log Transformation. To account for the heavy-tailed distribution of media coverage across goods, I apply a log transformation:

$$\tilde{n^M}_{i,t} = \log(1 + n^M_{i,t})$$

This reduces the influence of extreme values while preserving cross-good variation in media salience.

Normalization of Media Mentions Given Multiple Sources To ensure that media attention signals reflect genuine salience rather than differences in publication volume across outlets, I implement a normalization procedure at the source level. Specifically, for each news outlet s, I compute the share of articles mentioning good i in month t relative to the total number of articles published by that outlet in the same month:

$$\text{Share}^{(s)}_{i,t} = \frac{\text{Mentions of good } i \text{ in source } s \text{ at time } t}{\text{Total articles in source } s \text{ at time } t}$$

These shares are then averaged across all sources:

$$\text{MediaMentions}_{i,t} = \frac{1}{S} \sum_{s=1}^{S} \text{Share}^{(s)}_{i,t}$$

To address the heavy-tailed distribution of media coverage, I apply a log transformation:

$$\widetilde{\text{MediaMentions}}_{i,t} = \log\big(1 + \text{MediaMentions}_{i,t}\big)$$

I introduce source weights ψ_s based on audience size or reach. The weight ψ_s represents the proportion of the total audience across all outlets that consumes content from outlet s. Formally, if $\mathcal{R}_s$ denotes the estimated audience reach of source s, then

$$\psi_s = \frac{\mathcal{R}_s}{\sum_{j=1}^{S} \mathcal{R}_j}.$$

$$\text{MediaMentions}_{i,t} = \sum_{s=1}^{S} \psi_s \cdot \text{Share}_{i,t}^{(s)}.$$

This method ensures that the media attention index reflects *relative salience within each outlet*, avoids bias toward high-volume publishers, and appropriately accounts for the varying influence of different news sources based on their audience reach. It also allows for comparability across goods and time periods by focusing on *proportional coverage rather than raw mention counts.* Appendix C presents the time-series plots of normalized media mentions across all components.

Search Behavior

Search-based attention is proxied by the frequency of public information-seeking behavior related to specific goods. I estimate using Google Trends search indices for the same terms listed in Table 2.

Let $\mathcal{K}_i$ denote the set of search terms associated with good i. For each keyword, a time series is collected representing monthly search interest. Google Trends provides normalized search volume on a scale from 0 to 100.

$$\text{SearchIndex}_{i,t} = \text{GT}_{i,t}$$

where $\text{GT}_{i,t}$ denotes the Google Trends index for item i in month t. If multiple keywords are associated with a category, I take the average:

$$\text{SearchIndex}_{i,t} = \frac{1}{|\mathcal{K}_i|} \sum_{k \in \mathcal{K}_i} \text{GT}_{k,t}$$

Scale Alignment. Because Google Trends data is already scaled between 0 and 100, I use the index directly in the composite score.

Normalizations: Log-Transform or Smoothing. To address high volatility or outliers in Google Trends data, I (optionally) transform the raw search index using either a log transformation or a moving average smoother.

- **Log Transformation:**

$$\widetilde{\text{SearchIndex}}_{i,t} = \log(1 + \text{GT}_{i,t})$$

  This compresses large spikes in search volume while preserving relative variation.

- **(Optional) Moving Average (MA) Smoothing:**

$$\widetilde{\text{SearchIndex}}_{i,t} = \frac{1}{win} \sum_{j=0}^{win-1} \text{GT}_{i,t-j}$$

  where *win* is the window size (e.g., 3 months). This approach dampens month-to-month noise in consumer search behavior.

Appendix D presents the plots of the resulting component search indices.

4.2 On the Use of Price Changes $\pi_{i,t}$ in the PI.

It may appear circular to construct a perceived inflation using the same price changes that define the CPI. However, the innovation of the Perceived Inflation (PI) lies not in redefining inflation itself, but in capturing how individuals psychologically *weight* those price changes. Whereas the CPI uses fixed expenditure-based weights, the PI applies *attention-driven* *weights* that reflect behavioral responses to features such as salience, volatility, and media exposure. In this framework, the price change $\pi_{i,t}$ serves both as an economic quantity and as an attentional stimulus.

Measurement consistency: While measures of salience and volatility are derived from the inflation rate, media- and search-based indicators capture attention to *prices* rather than inflation per se. This asymmetry reflects the fact that households typically attend to salient nominal price changes and public "price talk," not to formally computed inflation aggregates. Combining the two is therefore appropriate for a perception-based index: perceived inflation is shaped jointly by objective variation in prices and by exposure to price narratives.

4.3 State-Space Measurement of Latent Attention and the Kalman Gain

The theoretical framework implies that perceived-inflation weights are determined by expenditure shares and the Kalman gain:

$$\omega_{i,t} = \frac{w_i^{\exp} K_{i,t}}{\sum_j w_j^{\exp} K_{j,t}}.$$

The Kalman gain is not directly observed. Instead of estimating it directly, I estimate a latent attention state that governs signal precision. This distinction is important: attention is not the Kalman gain itself. Rather, attention affects the noise of the price signal, and the Kalman gain follows mechanically from the signal-to-noise ratio.

Let $A_{i,t}$ denote latent attention to category i in month t. The attention state evolves according to

$$A_{i,t} = \rho A_{i,t-1} + u_{i,t}, \qquad u_{i,t} \sim \mathcal{N}(0, q). \tag{35}$$

The observed attention proxies are noisy measurements of this latent state. For each category i, month t, and proxy $j \in \{\text{salience}, \text{volatility}, \text{media}, \text{search}\}$, I write

$$z_{i,t}^{(j)} = \lambda_j A_{i,t} + \nu_{i,t}^{(j)}, \qquad \nu_{i,t}^{(j)} \sim \mathcal{N}(0, r_j), \tag{36}$$

where $z_{i,t}^{(j)}$ denotes the standardized value of proxy j. The salience loading is normalized to one, $\lambda_{\text{salience}} = 1$, which fixes the scale of the latent state.

Latent attention affects perceived inflation through signal precision. Specifically, the variance of the noisy price signal is assumed to satisfy

$$\log \sigma_{\epsilon,i,t}^2 = \alpha_0 - \alpha_1 A_{i,t}, \qquad \alpha_1 > 0. \tag{37}$$

Thus, higher attention lowers signal noise. Given the item-level inflation variance $\sigma_{\pi,i}^2$, the implied Kalman gain is

$$K_{i,t} = \frac{\sigma_{\pi,i}^2}{\sigma_{\pi,i}^2 + \sigma_{\epsilon,i,t}^2}. \tag{38}$$

Perceived inflation is then constructed as

$$\pi_t^{PI} = \sum_i \omega_{i,t} \pi_{i,t}, \qquad \omega_{i,t} = \frac{w_i^{\exp} K_{i,t}}{\sum_j w_j^{\exp} K_{j,t}}. \tag{39}$$

To discipline the mapping from latent attention to inflation perceptions, the baseline specification includes a survey-expectations anchor. Let χ_t denote the perceived inflation wedge,

$$\chi_t = \pi_t^{PI} - \pi_t^{CPI}. \tag{40}$$

The survey equation is

$$Survey_{t+1} = \beta_0 + \beta_{CPI}\pi_t^{CPI} + \beta_{Wedge}\chi_t + e_{t+1}, \qquad e_{t+1} \sim \mathcal{N}(0, \sigma_e^2). \tag{41}$$

The full system is estimated jointly by maximum likelihood using an extended Kalman filter and smoother. The measurement block identifies the latent attention state from salience, volatility, media mentions, and search behavior. The structural mapping from attention to signal precision then implies the Kalman gain, which in turn determines perceived-inflation weights. The survey equation provides an external behavioral anchor for the perception wedge.

4.4 Benchmark Specification and Robustness

The benchmark specification uses absolute price-change salience and normalized media mentions. Absolute salience is defined as $|\pi_{i,t}|$ and captures the idea that households are more likely to notice large price movements regardless of sign. Normalized media mentions adjust for differences in publication volume and therefore provide a cleaner measure of category-specific media salience.

The main specification is therefore denoted PI-AN-SSM, where AN indicates absolute salience for normalized media mentions, and SSM indicates the state-space measurement model. The New York Fed one-year-ahead inflation expectations series is used as the benchmark survey anchor. Michigan Survey mean and median expectations are used as robustness checks.

For comparison with the earlier reduced-form implementation, I also estimate variants using raw media mentions and deviation-from-average salience, $|\pi_{i,t} - \bar{\pi}_t|$. These alternatives produce very similar perceived-inflation paths, indicating that the main results are not driven by the media normalization or the precise salience measure.

4.5 Results

Table 3 presents the parameter estimates. The estimated coefficient on the PI wedge is negative in the benchmark specification: $\beta_{Wedge} = -1.3379$. Conditional on realized CPI inflation, a more positive PI wedge is associated with lower subsequent inflation expectations. To explore this deeper, I ran individual models for the CPI and PI. When entered separately, both CPI inflation and perceived inflation are positive and statistically significant predictors of survey expectations. The negative coefficient on the perceived-inflation wedge therefore does not imply that higher perceived inflation lowers expected inflation. Rather, it reflects the fact that CPI inflation contains a stronger signal for future inflation expectations than the attention-driven component of perceived inflation. Conditional on CPI inflation, the remaining variation captured by the wedge contributes negatively to the prediction of survey expectations.

The estimated persistence parameter is high, $\rho = 0.8089$, indicating that inflation attention evolves gradually over time. The estimated attention-to-precision elasticity is also positive and economically sizable, $\alpha_1 = 1.7256$, implying that higher attention lowers signal noise and mechanically increases the Kalman gain.

Table 3: Structural Parameter Estimates: Benchmark PI-AN-SSM Specification

Parameter	Interpretation	Estimate
ρ	Attention persistence	0.8089
q	Attention innovation variance	0.1517
α_0	Baseline signal-noise level	2.1711
α_1	Attention-to-precision elasticity	1.7256
β_0	Survey intercept	1.9988
β_{CPI}	CPI inflation coefficient	5.4135
β_{Wedge}	Perceived-inflation wedge coefficient	-1.3379

Notes: The benchmark specification corresponds to the PI-AN-SSM model using absolute inflation salience, normalized media attention, and the New York Fed Survey of Consumer Expectations as the survey anchor. Parameters are estimated by maximum likelihood using the extended Kalman filter. Measurement-loading coefficients are omitted because several loading parameters are weakly identified in the benchmark specification due to near-collinearity among attention proxies.

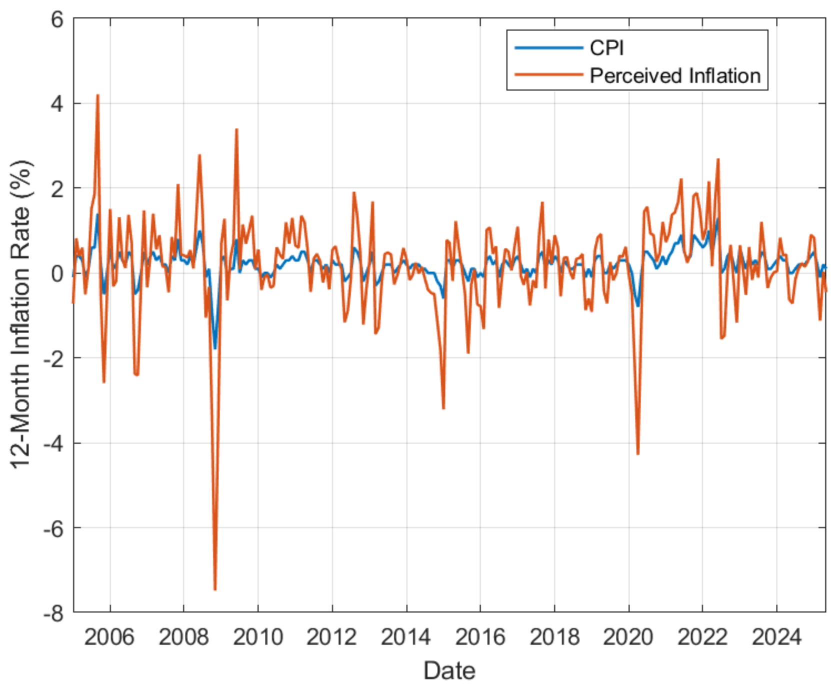

Figure 1: Headline CPI Inflation and the Benchmark State-Space Perceived Inflation Index (PI-AN-SSM)

Figure 1 compares the benchmark perceived-inflation index (PI-AN-SSM) with headline CPI inflation. The estimated perceived inflation series closely tracks aggregate inflation dynamics while exhibiting larger movements during periods of heightened macroeconomic stress. In particular, the perceived-inflation series displays amplified fluctuations during the Global Financial Crisis, the COVID-19 recession, and the 2021–22 inflation surge.

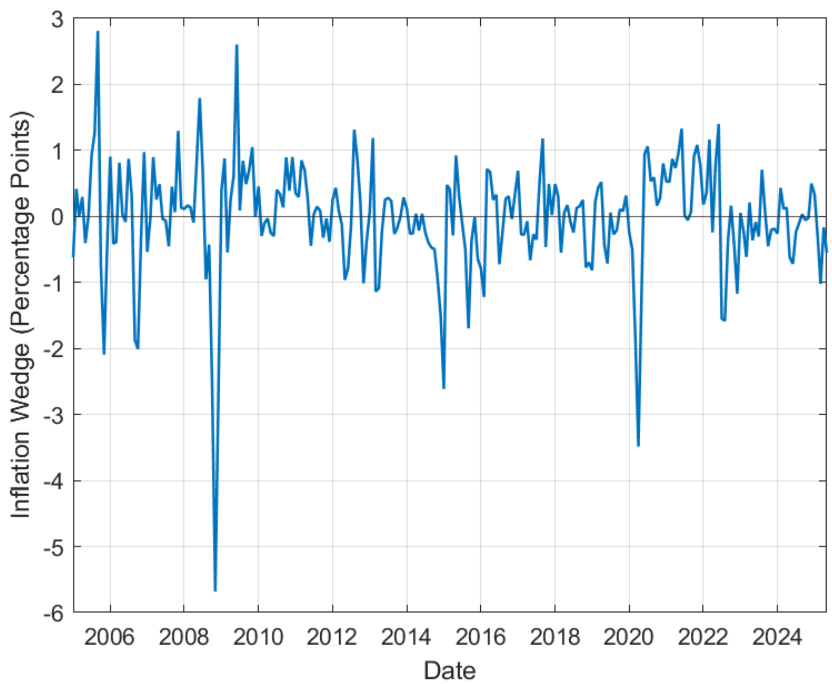

Figure 2: Perceived-Inflation Wedge

Figure 2 plots the perceived-inflation wedge,

$$\chi_t = \pi_t^{PI} - \pi_t^{CPI}.$$

The wedge fluctuates around zero over most of the sample but becomes strongly positive during periods of elevated inflation attention and strongly negative during severe downturns. The largest deviations occur during the financial crisis, the pandemic recession, and the post-pandemic inflation episode, indicating that household inflation perceptions diverge most sharply from official CPI inflation during periods of macroeconomic disruption.

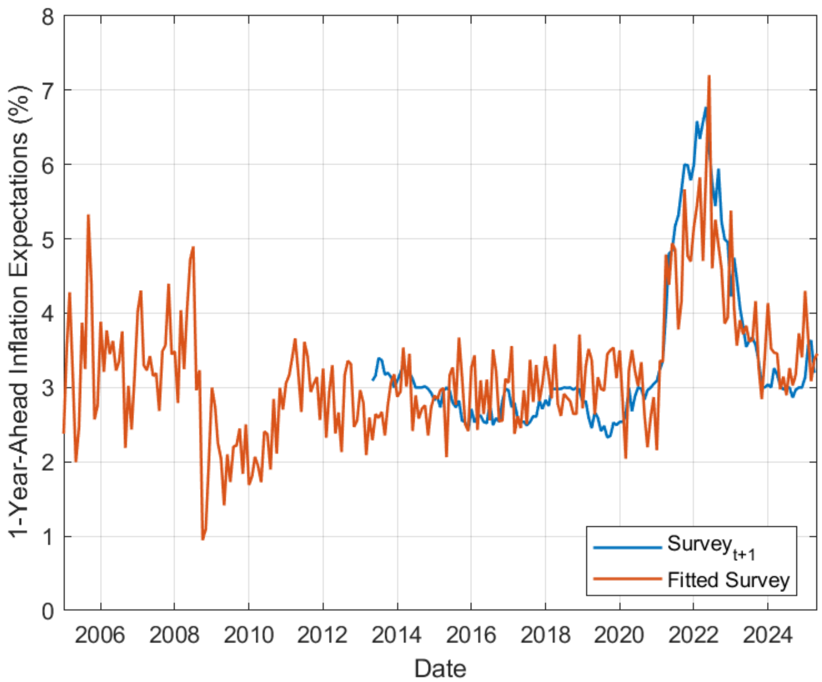

Figure 3: Survey Anchor Fit: New York Fed Inflation Expectations

Figure 3 compares the model-implied survey anchor with observed one-year-ahead inflation expectations from the New York Fed Survey of Consumer Expectations. The state-space model reproduces both medium-run movements and major turning points in survey expectations, including the sharp increase in expected inflation during 2021–22 and the subsequent decline.

5 Application: Perceived Inflation in a Heterogeneous Household Model

The primary contribution of this paper is the construction and estimation of a behaviorally grounded measure of perceived inflation. An important remaining question is whether this measure has useful macroeconomic applications. To illustrate one potential use, this section embeds the estimated perceived-inflation wedge into a standard two-asset heterogeneous-household framework. The objective is to demonstrate how empirically observed variation in perceived inflation can affect saving, consumption, and liquidity outcomes when households base intertemporal decisions on perceived rather than official inflation.

5.1 Economic Environment

Time is discrete. The economy consists of a continuum of households facing idiosyncratic income risk and incomplete markets. Households have access to two assets: A liquid asset $a_{i,t}$ earning return R_t^A, and an illiquid asset $k_{i,t}$ earning return R_t^K, subject to adjustment costs.

Households solve:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{C_{i,t}^{1-\sigma} - 1}{1 - \sigma}, \tag{42}$$

subject to

$$C_{i,t} + a_{i,t+1} + k_{i,t+1} + \Psi(k_{i,t+1}, k_{i,t}) = R_t^A a_{i,t} + R_t^K k_{i,t} + Y_{i,t}, \tag{43}$$

and borrowing constraint

$$a_{i,t+1} \geq \underline{a}. \tag{44}$$

5.2 Perceived Inflation and the Attention-Based Wedge

In the model, households do not observe CPI inflation directly. Instead, they form beliefs about inflation based on attention-weighted signals of individual price changes, as developed in Section 3.

Perceived inflation is given by:

$$\pi_t^{PI} = \sum_i \omega_{i,t} \pi_{i,t},$$

where the weights $\omega_{i,t}$ depend on endogenous attention and information precision through the Kalman gain:

$$\omega_{i,t} = \frac{w_i^{exp} K_{i,t}}{\sum_j w_j^{exp} K_{j,t}}.$$

This implies that perceived inflation differs from CPI inflation:

$$\pi_t^{PI} = \pi_t^{CPI} + \chi_t,$$

where the wedge χ_t reflects distortions in attention across goods.

Importantly, χ_t is not exogenous. It is an equilibrium object that arises from: (i) optimal signal-acquisition effort, (ii) variation in signal precision, and (iii) the external information environment (media and search activity).

5.3 Perceived Inflation in Intertemporal Decisions

Households make intertemporal decisions based on perceived inflation rather than official CPI inflation. In particular, the real return relevant for consumption decisions depends on perceived inflation:

$$r_t^{real} = r_t - \mathbb{E}_t \pi_{t+1}^{PI}.$$

This implies that the Euler equation governing consumption choices is given by

$$C_{i,t}^{-\sigma} = \beta \, \mathbb{E}_t \left[C_{i,t+1}^{-\sigma} \left(1 + r_t - \pi_{t+1}^{PI}\right) \right].$$

Substituting $\pi_{t+1}^{PI} = \pi_{t+1}^{CPI} + \chi_{t+1}$ shows that the perception wedge directly distorts households' perception of real interest rates and therefore affects saving and consumption decisions.

5.4 Dynamics of the Perception Wedge

The perception wedge χ_t is microfounded through attention and information frictions, but solving the full dynamic signal-acquisition problem inside the heterogeneous-agent model is computationally intractable. I therefore use a parsimonious reduced-form process for the empirically estimated wedge:

$$\pi_t^{PI} = \pi_t^{CPI} + \chi_t,$$

where

$$\chi_t = \rho_\chi \chi_{t-1} + \gamma_\chi \varepsilon_t^{oil} + \eta_t.$$

This specification preserves the structural interpretation of χ_t as an attention-driven distortion in perceived inflation while allowing the household model to incorporate empirically disciplined wedge dynamics. The oil supply news shock ε_t^{oil} (source: Känzig (2021)) provides an external source of inflationary news, while ρ_χ captures the persistence of perceived-inflation distortions.

Interpretation of the Wedge as an Attention Distortion The perception wedge χ_t can be interpreted as a distortion in the mapping between observed prices and perceived aggregate inflation. In the absence of attention frictions, all goods would receive informational weights proportional to their expenditure shares, implying $\chi_t = 0$.

With limited attention, households overweight salient, volatile, or highly publicized goods, leading to systematic deviations from CPI inflation. As a result, χ_t reflects the endogenous misallocation of attention across goods, rather than an exogenous belief shock.

The state-dependent process above distinguishes the present framework from models that introduce exogenous inflation belief shocks. Here, belief distortions are disciplined by observable attention data and by the estimated response of the perceived inflation wedge to identified oil supply news shocks. Thus, the wedge is not merely a reduced-form disturbance; it is a behaviorally grounded object whose quantitative relevance depends on the attention regime.

5.5 Calibration and Scenario Selection

I calibrate the household side of the model to match key moments of the wealth distribution and consumption behavior that are standard in the heterogeneous-agent literature. Specifically, I target four moments: the share of hand-to-mouth (HtM) households, two distributional measures of wealth in the economy (mean and median wealth), and the marginal propensity to consume (MPC).

To assess the robustness of the calibration and isolate the role of different empirical moments, I consider several alternative calibration specifications that vary the set of targeted moments. The most comprehensive specification targets all four moments simultaneously. A second specification focuses on matching the wealth distribution by targeting the HtM share together with mean and median wealth, while leaving the MPC unconstrained. In contrast, the third—and preferred—specification emphasizes liquidity and consumption behavior by targeting the HtM share, median wealth, and the MPC, abstracting from the mean wealth moment. Finally, the most parsimonious specification targets only the HtM share and the MPC, providing a minimal benchmark for capturing household liquidity and consumption responses.

5.6 Baseline Calibration

These alternative calibration exercises allow us to evaluate which empirical moments are most informative for the mechanisms central to the analysis. I adopt the liquidity-focused specification as the baseline because it provides the best balance between matching household liquidity, consumption responsiveness, and the empirically relevant distribution of wealth. Relative to the wealth-focused alternatives, the liquidity-centric specification more effectively captures hand-to-mouth behavior and MPC heterogeneity, both of which are central to understanding how perceived inflation distortions transmit into household consumption decisions.

Table 4 reports the calibrated household-side parameters. The estimated monthly discount factor ($\beta = 0.9975$) implies a relatively patient household sector, while the CRRA coefficient ($\sigma = 4.45$) indicates substantial precautionary saving motives. The borrowing limit ($\underline{a} = -1.02$) permits limited short-term liquidity smoothing, while the adjustment

Table 4: Baseline Household Calibration

Parameter	Description	Value
β	Monthly discount factor	0.9975
σ	CRRA coefficient	4.4531
$\underline{a}$	Borrowing limit	-1.0195
ϕ	Illiquid adjustment cost curvature	0.00273
F	Fixed adjustment cost	0.00082
r^A	Liquid real return	0.00106
r^K	Illiquid real return	0.00321

Notes: Parameter values correspond to the preferred LIQ3 calibration, which targets the hand-to-mouth share, median wealth-to-income ratio, and quarterly MPC.

cost parameters (ϕ and F) imply moderate frictions in reallocating between liquid and illiquid assets.

Table 5: Estimated Perception-Wedge Process

Parameter	Description	Estimate
ρ_χ	Wedge persistence	0.453
γ_χ	Oil-shock sensitivity	0.0118
σ_η	Innovation volatility	0.144

Notes: The perception wedge follows $\chi_t = \rho_\chi \chi_{t-1} + \gamma_\chi \varepsilon_t^{oil} + \eta_t$. The shock ε_t^{oil} is the standardized Känzig oil supply news shock. The process is estimated using the benchmark state-space perceived inflation wedge.

Table 5 reports the estimated reduced-form process governing the benchmark state-space perception wedge. The results indicate a moderate degree of persistence in perceived inflation distortions, with an autoregressive coefficient of 0.453. The estimated response of the wedge to oil supply news shocks is positive but relatively small, with a sensitivity parameter of 0.0118.

Table 6 reports household outcomes across alternative perceived-inflation wedge states. The CPI-only column corresponds to the benchmark economy in which perceived inflation

Table 6: Calibration Fit and Wedge Counterfactuals

	Target	CPI	q10	q25	q50	q75	q90
HtM share	0.14	0.133	0.123	0.137	0.115	0.098	0.123
Mean wealth / income	–	19.68	20.37	19.95	19.57	19.16	18.80
Median wealth / income	18.0	16.93	17.46	17.11	16.85	16.67	16.42
Quarterly MPC	0.20	0.377	0.476	0.424	0.380	0.349	0.359

Notes: The CPI column reports the benchmark economy in which the perceived-inflation wedge is set equal to zero, $\chi_t = 0$. The remaining columns evaluate household outcomes at different points of the empirical wedge distribution estimated from the benchmark state-space model. Specifically, q10, q25, q50, q75, and q90 correspond to the 10th, 25th, 50th, 75th, and 90th percentiles of the estimated perceived-inflation wedge distribution. Negative wedge realizations imply perceived inflation below official CPI inflation, while positive realizations imply perceived inflation above official CPI inflation.

equals official inflation. The remaining columns evaluate household behavior at different points of the empirical wedge distribution estimated from the state-space model. The results indicate that perceived-inflation distortions have meaningful effects on household balance sheets and consumption behavior. Variation in the wedge changes the fraction of hand-to-mouth households, wealth accumulation, and marginal propensities to consume, even though official inflation remains unchanged.

5.7 Distributional Effects of Perceived Inflation

While Table 6 reports aggregate household outcomes across alternative perceived-inflation wedge realizations, the effects are not uniform across households. Because perceived inflation alters households' perceived real returns, the resulting changes in saving and consumption behavior depend on their position in the wealth distribution. To examine these distributional implications, this exercise examines the model-implied responses across wealth quartiles.

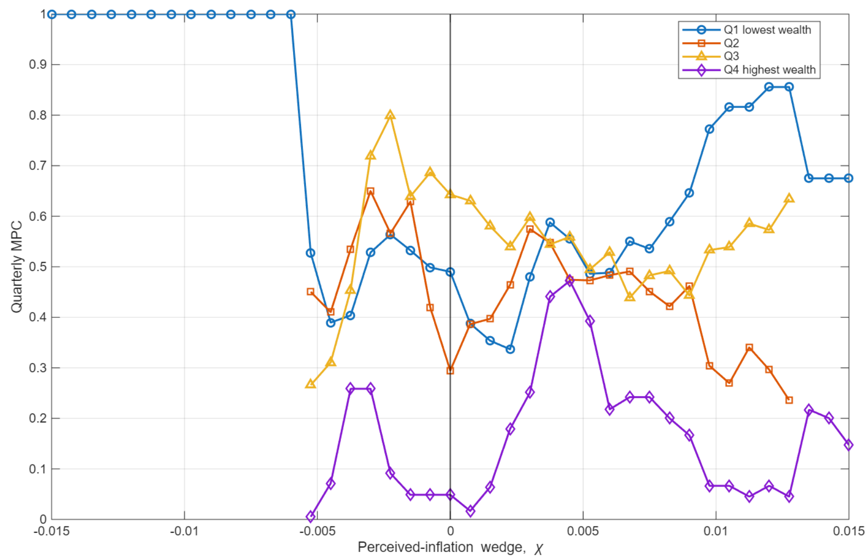

Figure 4: Quarterly MPC Responses Across Wealth Quartiles

Figure 4 shows substantial heterogeneity in consumption responses across the wealth distribution. The lowest-wealth households exhibit the largest marginal propensities to consume throughout the range of perceived-inflation wedges considered, while the highest-wealth households display substantially lower MPCs. A plausible interpretation is the presence of liquidity constraints: households with limited wealth are more exposed to changes in perceived real returns and therefore adjust consumption more aggressively when inflation perceptions change.

Interestingly, for liquidity-constrained households, MPC responses are noticeably larger for positive wedges than for negative wedges. This asymmetry suggests that these households respond more strongly when perceived inflation rises above official inflation, consistent with a heightened sensitivity to declines in perceived real returns.

Figure 5 reports the fraction of households that behave as hand-to-mouth consumers

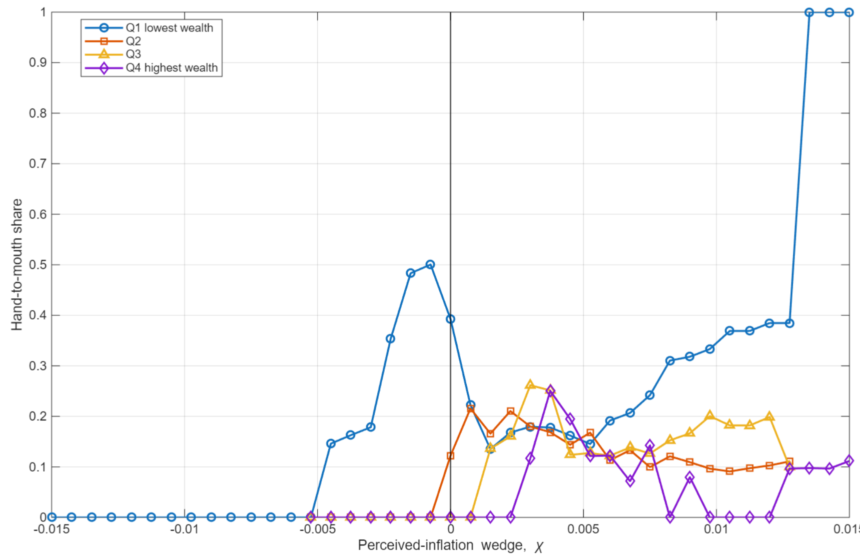

Figure 5: Hand-to-Mouth Responses Across Wealth Quartiles

across wealth quartiles. The perceived-inflation distortions primarily affect liquidity positions among households with limited financial buffers, while wealthier households remain somewhat insulated from these changes. That is, the effects of the perceived-inflation wedge are more concentrated among households at the lower end of the wealth distribution. Increases in the wedge are associated with noticeably larger movements in hand-to-mouth status among poorer households, whereas the highest-wealth households exhibit smaller variation. Finally, positive wedge realizations generate a pronounced increase in hand-to-mouth status among the poorest households, whereas the corresponding response becomes progressively weaker as wealth increases.

6 Conclusion

This paper develops a behavioral measure of inflation that captures how households perceive price changes. I introduce a Perceived Inflation (PI) index that reweights price changes based on salience, volatility, media coverage, and search behavior, and show how these attention-based weights arise endogenously within a microfounded framework. In the model, households optimally allocate scarce attention across goods, implying that more salient, volatile, or information-rich prices receive disproportionately greater weight in perceived inflation. This generates a time-varying wedge between official and perceived inflation that emerges structurally from attention allocation and information frictions.

Empirically, perceived inflation differs systematically from CPI inflation. Although the two series move closely together over time, perceived inflation exhibits substantially greater short-run volatility. The estimated perceived-inflation wedge is weakly persistent and mean-reverting, indicating that attention-driven distortions are systematic but not permanent. The resulting measure closely tracks survey-based inflation expectations and provides information about household beliefs beyond realized inflation alone.

To assess its structural implications, I embed perceived inflation into a perception-augmented heterogeneous-agent model. In this framework, households make intertemporal decisions using perceived rather than official inflation, so deviations between perceived and realized inflation affect perceived real returns, saving behavior, and consumption choices. The model implies that empirically observed variation in the perceived-inflation wedge has economically meaningful consequences for household balance sheets and consumption outcomes. Changes in perceived inflation influence wealth accumulation, liquidity positions, and marginal propensities to consume, with the largest responses concentrated among low-wealth and liquidity-constrained households.

References

Bates, J. M. and Gabor, A. (1986). Price perception in creeping inflation: Report on an enquiry. *Journal of Economic Psychology*, 7(3):291–314.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.

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

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

D'Acunto, F., Malmendier, U., and Weber, M. (2023). What do the data tell us about inflation expectations? In Coibion, O., Gorodnichenko, Y., and Weber, M., editors, *Handbook of Economic Expectations*, pages 133–161. Academic Press.

Detmeister, A., Lebow, D., and Peneva, E. (2016). Inflation perceptions and inflation expectations. *FEDS Notes*. Washington: Board of Governors of the Federal Reserve System.

Gabaix, X. (2014). A sparsity-based model of bounded rationality *. *The Quarterly Journal of Economics*, 129(4):1661–1710.

Georganas, S., Healy, P. J., and Li, N. (2014). Frequency bias in consumers' perceptions of inflation: An experimental study. *European Economic Review*, 67:144–158.

Jonung, L. (1981). Perceived and expected rates of inflation in sweden. *The American Economic Review*, 71(5):961–968.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4):1092–1125.

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

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

Pfäuti, O. (2026). The inflation attention threshold and inflation surges. *American Economic Review*. Conditionally accepted; previous version available as arXiv:2308.09480.

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

Weber, M., D'Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, 36(3):157–184.

Appendix

A Proofs

Appendix A Proof of Proposition 1

Fix a good i and a period t, and suppress the indices (i,t) for notational simplicity. Let $\sigma_\pi^2 \equiv \sigma_{\pi,i}^2$, $w \equiv w_i^{\exp}$, $\ell \equiv \ell_{i,t}$, and $\bar{\sigma}^2 \equiv \bar{\sigma}_i^2$. The per-good contribution to the period-t objective is

$$J(\ell;\sigma_\pi^2,w) = \frac{\gamma}{2}w^2(1-K(\ell))\sigma_\pi^2 + \frac{\lambda}{2}\ell^2,$$

where

$$\sigma_\varepsilon^2(\ell) = \frac{\bar{\sigma}^2}{1+\phi\ell}, \qquad K(\ell) = \frac{\sigma_\pi^2}{\sigma_\pi^2+\sigma_\varepsilon^2(\ell)}, \qquad 1-K(\ell) = \frac{\sigma_\varepsilon^2(\ell)}{\sigma_\pi^2+\sigma_\varepsilon^2(\ell)}.$$

The first-order condition (FOC) for the optimal signal-acquisition effort $\ell^*(\sigma_\pi^2,w)$ is

$$F(\ell,\sigma_\pi^2,w) \equiv \frac{\partial J}{\partial \ell} = \frac{\gamma}{2}w^2\sigma_\pi^2\frac{\partial}{\partial \ell}(1-K(\ell)) + \lambda\ell = 0. \tag{45}$$

We analyze $\ell^*(\sigma_\pi^2,w)$ using the implicit function theorem.

Step 1: Sign of $\frac{\partial}{\partial \ell}(1-K(\ell))$. First compute $\frac{d\sigma_\varepsilon^2}{d\ell}$:

$$\frac{d\sigma_\varepsilon^2}{d\ell} = -\frac{\phi\bar{\sigma}^2}{(1+\phi\ell)^2} < 0.$$

Treating $1-K$ as a function of σ_ε^2,

$$1-K = \frac{\sigma_\varepsilon^2}{\sigma_\pi^2+\sigma_\varepsilon^2}, \qquad \frac{\partial(1-K)}{\partial\sigma_\varepsilon^2} = \frac{\sigma_\pi^2}{\left(\sigma_\pi^2+\sigma_\varepsilon^2\right)^2} > 0.$$

By the chain rule,

$$\frac{\partial}{\partial \ell}(1-K(\ell)) = \frac{\partial(1-K)}{\partial \sigma_\varepsilon^2} \cdot \frac{d\sigma_\varepsilon^2}{d\ell} = \frac{\sigma_\pi^2}{\left(\sigma_\pi^2 + \sigma_\varepsilon^2(\ell)\right)^2} \left(-\frac{\phi \bar{\sigma}^2}{(1+\phi\ell)^2} \right) < 0. \tag{46}$$

Define

$$B(\ell, \sigma_\pi^2) \equiv \frac{\partial}{\partial \ell}(1-K(\ell)),$$

so (46) implies

$$B(\ell, \sigma_\pi^2) < 0 \quad \text{for all } \ell \geq 0,\ \sigma_\pi^2 > 0. \tag{47}$$

Using this notation, the FOC (45) becomes

$$F(\ell, \sigma_\pi^2, w) = \frac{\gamma}{2} w^2 \sigma_\pi^2 B(\ell, \sigma_\pi^2) + \lambda \ell = 0. \tag{48}$$

Step 2: Second-order condition and $F_\ell > 0$. Differentiating F with respect to ℓ gives

$$F_\ell(\ell, \sigma_\pi^2, w) = \frac{\gamma}{2} w^2 \sigma_\pi^2 \frac{\partial B}{\partial \ell}(\ell, \sigma_\pi^2) + \lambda.$$

We assume that the second-order condition holds at the optimum, so that

$$F_\ell(\ell^*, \sigma_\pi^2, w) > 0.$$

This condition is satisfied, for example, when the quadratic cost term is sufficiently strong relative to the curvature of the perception-loss term. Under this regularity condition, the optimum is locally unique and the implicit function theorem applies.

Step 3: Comparative statics with respect to w. Differentiating (48) with respect to w gives

$$F_w(\ell, \sigma_\pi^2, w) = \gamma w \sigma_\pi^2 B(\ell, \sigma_\pi^2).$$

Using $\gamma > 0$, $\sigma_\pi^2 > 0$, $w > 0$, and $B(\ell, \sigma_\pi^2) < 0$ from (47), we have

$$F_w(\ell, \sigma_\pi^2, w) < 0.$$

By the implicit function theorem,

$$\frac{\partial \ell^*}{\partial w} = -\frac{F_w}{F_\ell}\bigg|_{\ell=\ell^*(\sigma_\pi^2, w)} > 0. \tag{49}$$

Thus optimal signal-acquisition effort is strictly increasing in the expenditure weight $w_i^{\exp}$.

Step 4: Comparative statics with respect to σ_π^2. For the volatility derivative it is convenient to work with $\sigma_\pi^2 B$. For any fixed ℓ,

$$B(\ell, \sigma_\pi^2) = \frac{\sigma_\pi^2}{\left(\sigma_\pi^2 + \sigma_\varepsilon^2(\ell)\right)^2} \frac{d\sigma_\varepsilon^2}{d\ell}.$$

Let $x \equiv \sigma_\pi^2 > 0$ and $c \equiv \sigma_\varepsilon^2(\ell) > 0$, which does not depend on x, and let $d \equiv -\frac{d\sigma_\varepsilon^2}{d\ell} > 0$. Then

$$B(\ell, \sigma_\pi^2) = -\frac{d\,x}{(x+c)^2}, \qquad xB(\ell, \sigma_\pi^2) = -\frac{d\,x^2}{(x+c)^2}.$$

Define

$$G(x) \equiv xB(\ell, \sigma_\pi^2) = -d\,\frac{x^2}{(x+c)^2}.$$

Differentiating,

$$G'(x) = -d\,\frac{2x(x+c)^2 - x^2 \cdot 2(x+c)}{(x+c)^4} = -d\,\frac{2xc}{(x+c)^3} < 0 \qquad \text{for all } x > 0.$$

Hence

$$\frac{\partial}{\partial \sigma_\pi^2}\big(\sigma_\pi^2 B(\ell, \sigma_\pi^2)\big) < 0 \quad \text{for all } \ell \geq 0,\ \sigma_\pi^2 > 0. \tag{50}$$

Differentiating (48) with respect to σ_π^2 therefore gives

$$F_{\sigma_\pi^2}(\ell, \sigma_\pi^2, w) = \frac{\gamma}{2} w^2 \frac{\partial}{\partial \sigma_\pi^2}\big(\sigma_\pi^2 B(\ell, \sigma_\pi^2)\big) < 0.$$

Again by the implicit function theorem,

$$\frac{\partial \ell^*}{\partial \sigma_\pi^2} = -\frac{F_{\sigma_\pi^2}}{F_\ell}\bigg|_{\ell=\ell^*(\sigma_\pi^2, w)} > 0. \tag{51}$$

Equations (49) and (51) establish that optimal signal-acquisition effort is strictly increasing in both $\sigma_{\pi,i}^2$ and $w_i^{\exp}$:

$$\frac{\partial \ell_{i,t}}{\partial \sigma_{\pi,i}^2} > 0, \qquad \frac{\partial \ell_{i,t}}{\partial w_i^{\exp}} > 0.$$

Step 5: Implications for perceived-inflation weights. Finally, recall that

$$\omega_{i,t} = \frac{w_i^{\exp} K_{i,t}}{\int_0^1 w_j^{\exp} K_{j,t}\, dj}, \qquad K_{i,t} = \frac{\sigma_{\pi,i}^2}{\sigma_{\pi,i}^2 + \sigma_{\varepsilon,i}^2(\ell_{i,t})}.$$

Because $\sigma_{\varepsilon,i}^2(\ell)$ is strictly decreasing in ℓ, the Kalman gain $K_{i,t}$ is strictly increasing in $\ell_{i,t}$. Therefore, goods with higher volatility $\sigma_{\pi,i}^2$ and larger expenditure share $w_i^{\exp}$ receive greater signal-acquisition effort, higher Kalman gains, and larger perceived-inflation weights $\omega_{i,t}$. $\square$

Appendix B Proof of Proposition 2 (Media Exposure, Search, and Perceived Inflation)

Fix good i and period t and suppress indices (i,t) for brevity. Let $\mathcal{M}$ denote media exposure and $\mathcal{S}$ denote search intensity. The signal variance is

$$\sigma_\varepsilon^2(\ell, \mathcal{M}, \mathcal{S}) = \frac{\bar{\sigma}^2}{1 + \phi\ell + \phi_M \mathcal{M} + \phi_S \mathcal{S}}, \qquad \bar{\sigma}^2 > 0,\ \phi, \phi_M, \phi_S > 0,$$

and the Kalman gain is

$$K = \frac{\sigma_\pi^2}{\sigma_\pi^2 + \sigma_\varepsilon^2(\ell, \mathcal{M}, \mathcal{S})}, \qquad \sigma_\pi^2 > 0.$$

Step 1: Media exposure and search intensity raise the Kalman gain. First note that

$$\frac{\partial \sigma_\varepsilon^2}{\partial \mathcal{M}} = -\frac{\phi_M \bar{\sigma}^2}{\left(1 + \phi\ell + \phi_M \mathcal{M} + \phi_S \mathcal{S}\right)^2} < 0,$$

and

$$\frac{\partial \sigma_\varepsilon^2}{\partial \mathcal{S}} = -\frac{\phi_S \bar{\sigma}^2}{\left(1 + \phi\ell + \phi_M \mathcal{M} + \phi_S \mathcal{S}\right)^2} < 0.$$

Treating K as a function of σ_ε^2,

$$\frac{\partial K}{\partial \sigma_\varepsilon^2} = -\frac{\sigma_\pi^2}{\left(\sigma_\pi^2 + \sigma_\varepsilon^2\right)^2} < 0.$$

By the chain rule,

$$\frac{\partial K}{\partial \mathcal{M}} = \frac{\partial K}{\partial \sigma_\varepsilon^2} \cdot \frac{\partial \sigma_\varepsilon^2}{\partial \mathcal{M}} > 0, \qquad \frac{\partial K}{\partial \mathcal{S}} = \frac{\partial K}{\partial \sigma_\varepsilon^2} \cdot \frac{\partial \sigma_\varepsilon^2}{\partial \mathcal{S}} > 0.$$

Thus

$$\frac{\partial K_{i,t}}{\partial \mathcal{M}_{i,t}} > 0, \qquad \frac{\partial K_{i,t}}{\partial \mathcal{S}_{i,t}} > 0.$$

Step 2: Effects on perceived-inflation weights. Perceived-inflation weights are

$$\omega_{i,t} = \frac{w_i^{\exp} K_{i,t}}{\int_0^1 w_j^{\exp} K_{j,t}\, dj}.$$

Holding $(K_{j,t})_{j \neq i}$ fixed, the map

$$K_{i,t} \mapsto \frac{w_i^{\exp} K_{i,t}}{\int_0^1 w_j^{\exp} K_{j,t}\, dj}$$

is strictly increasing in $K_{i,t}$ for $w_i^{\exp} > 0$. Since media exposure and search intensity increase $K_{i,t}$, they also increase the perceived-inflation weight $\omega_{i,t}$. Therefore, goods receiving greater media coverage or search activity receive larger perceived-inflation weights and disproportionately influence perceived inflation. □

Appendix C Obtaining Media Mentions

Data Collection and Processing

The analysis relies on a corpus of XML-based text files representing news articles, reports, and related publications organized by consumption categories corresponding to the Consumer Price Index (CPI). The raw data were stored in a hierarchical directory structure, with each folder representing a CPI category (e.g., *Shelter*, *Fuel and Utilities*, *Public Transportation*). Within each category, individual XML files contained article-level metadata and textual content.

Data Extraction

To extract relevant information, each XML file was parsed using Python's `lxml` and `BeautifulSoup` libraries. For each document, key metadata fields were retrieved—such as a unique identifier, publication title, date, publisher, country of origin, and full text content. HTML tags embedded in text fields were removed to ensure clean textual analysis. The extraction process included error handling to bypass corrupted or incomplete files.

A simplified example of the extracted structure includes:

- **Document ID:** unique article identifier
- **Title:** article headline or publication title
- **Date:** publication date (converted to YYYY-MM format)
- **Publisher and Country:** source organization and geographic location
- **Text:** cleaned article body

Only documents published in the **United States** were retained for further analysis to maintain geographic consistency.

Keyword Filtering and Term Counting

Each CPI category was assigned a predefined list of search terms representing goods or services relevant to that category (for example, the *Fuel and Utilities* category included terms such as “gas prices” or “utility bills”). For each document, the code computed the frequency of these terms within both the title and body text. Regular expressions were used to perform case-insensitive searches, and total monthly counts were aggregated to the category level.

A typical operation for a given document i involved:

$$\text{count}_{it} = \sum_{k=1}^{K} \text{Occurrences of term } k \text{ in document } i$$

These counts were then aggregated by month and publisher, yielding two metrics per month:

1. **Raw total counts** – the total number of mentions of relevant terms in all documents during the month.

2. **Normalized counts** – term frequencies adjusted by the number of unique publishers to mitigate the influence of large outlets.

Parallel Processing and Efficiency

Given the size of the corpus, data extraction and processing were parallelized using Python’s `multiprocessing` module. Files within each category were divided into batches, each processed concurrently across available CPU cores. This approach significantly reduced computation time when parsing and analyzing tens of thousands of XML files per CPI category.

Output and Data Structure

For each CPI category, two output datasets were generated and saved in `.csv` format:

- **Publisher-level counts:** the number of relevant term mentions per publisher per month.
- **Monthly normalized counts:** aggregated and adjusted monthly frequency metrics across all publishers.

These datasets form the foundation for subsequent measures of media attention and perceived inflation. The resulting monthly time series span from 2004 to 2025 and are used in constructing attention-weighted inflation perception indices.

Reproducibility Note

All data processing was performed in `Python 3.12` within an `AWS SageMaker` environment. The key libraries included `pandas` (for data manipulation), `lxml` and `BeautifulSoup` (for XML parsing and text cleaning), and `multiprocessing` (for parallel execution). The script automatically detected available CPU cores and utilized all but one for concurrent batch processing. Intermediate outputs were saved to structured directories, enabling full reproducibility of the extraction, cleaning, and aggregation steps.

Event Study

We will conduct event studies around periods of:

- Large, salient price shocks (e.g., oil price spikes, food crises)
- Major media coverage events linked to prices (e.g., record gasoline prices)

We will examine whether major events identified over time align with key inflationary (or deflationary) episodes, over the sample period, for the corresponding categories of the CPI (see Table 7).

Table 7: Major U.S. Price Shocks Since 2000 Across Key Goods

Period	**Event & Cause**	**Affected Prices & Timing**
2005	Oil price surge due to geopolitical tensions and Hurricane Katrina	Brent crude surpassed \$60/barrel; U.S. gasoline prices rose sharply. *Aug 2005: Brent crude ¿ \$60/bbl*
2003–2008	Oil price surge due to global demand, OPEC constraints, Middle East risk	Crude oil rose from $<$ \$30 to $\sim$ \$147/barrel; U.S. gasoline peaked around \$4.10/gal. *July 2008: Oil reached \$147/bbl*
2006–2008	Global food crisis (droughts, ethanol mandates, speculation)	Rice $\uparrow >$ 230%, wheat $\uparrow$ 136%, corn $\uparrow$ 125% globally. *April 2008: Rice exceeded \$1,000/ton*
Late 2008	Global demand crash during Great Recession	Oil fell to $\sim$ \$32/barrel; gasoline dropped below \$2/gal. *Dec 2008: Oil at \$32/bbl*
2010–2011	Global food price surge (poor harvests, rising oil/fertilizer costs, biofuel demand, speculation, export restrictions)	FAO Food Price Index rose from 188 in 2010 to a peak of 238 points; cereals alone soared 35% in the year to February 2011. *Feb 2011: FFPI nominal peak 237.9; cereals surged 35% YoY*
2014–2016	U.S. shale boom $\rightarrow$ oil oversupply	Oil dropped from \$100 to $<$ \$30/barrel. *Jan 2016: Oil below \$30/bbl*
2020–2021	COVID-19 supply shocks + demand surge (DIY, housing, fiscal stimulus)	Lumber prices surged to \$1,686/MBF; food and energy prices rose sharply. *May 2021: Lumber peak*
2022–2023	Russia–Ukraine war + fertilizer, fuel, and export disruptions	Wheat, sunflower oil, and fertilizer prices spiked; global food price index hit record highs in early 2022. Severe impact in import-dependent nations. *Mar 2022: FAO Food Price Index peaked at record level*

Continued on next page

Table 7 – continued from previous page

Period	Event & Cause	Affected Prices & Timing
2023–2025	Disinflation as supply chains normalize; Fed tightening	Food CPI stabilized near ∼ 3% YoY by mid-2025, from a high of 8.5% May 2022; energy prices moderated - the Global Energy Index inflation YoY fell from 104.7% August 2022 to close to zero (at times negative) since January 2023 *July 2025: Gasoline around $3.50/gal*

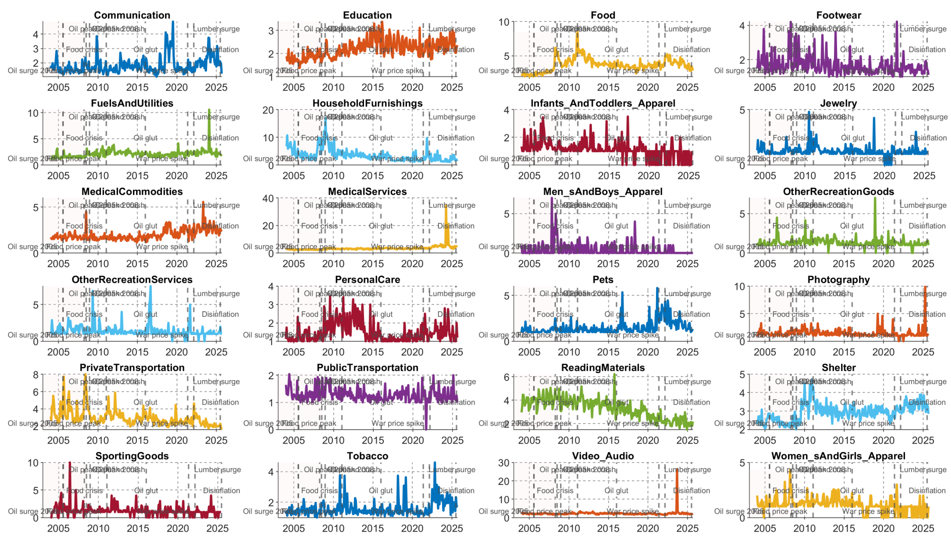

Figure 6: Normalized Media mentions by component with global event markers and oil-window shading.

Joint Reading of Table 7 and Figure 6 Table 7 lists the principal price shocks over the sample period, while Figure 6 displays normalized media-mention intensity across individual CPI categories. The two exhibit a close temporal correspondence. Specifically, categories such as *Fuel and Energy*, *Food and Nonalcoholic Beverages*, and *Transport* display pronounced spikes in media attention coinciding with the 2003–2008 oil and food crises and again during the 2010–2011 commodity surge. A further broad-based uptick is visible in *Food*, *Housing and Utilities*, and *Transport* around 2020–2021, matching the pandemic-induced disruptions reported in Table 6. The 2022 war-related commodity shock produces a final synchronized peak across *Energy*, *Food*, and *Durables*.

By contrast, the closing segment of the figure—corresponding to the 2023–2025 disinflation period—shows a downward trajectory in most categories, particularly *Energy*, *Food*, and *Transport*, reflecting the easing of global supply constraints documented in Table 7.

Overall, the co-movement of these series confirms that the intensity of media coverage rises sharply in tandem with major price shocks and subsides as inflation pressures abate. The alignment substantiates the interpretation that media salience transmits objective cost shocks into perception shocks, with attention contracting during disinflationary phases.

Appendix D Computation of Relative Index Changes Across Crises

For each of the abovementioned categories (*Food*, *Energy*, *Private/Public Transport*, *Shelter*, *Durables*) plus an aggregated *AllOther* median index, the code computes the percentage deviation from a buffered pre-period extremum. During crisis periods, the statistic is

$$\text{Spike} = \frac{\max(\text{in-period})}{\min(\text{pre-period})} - 1,$$

while for the disinflation phase it reverses sign:

$$\text{Spike} = \frac{\min(\text{in-period})}{\max(\text{pre-period})} - 1.$$

All results are expressed in percent.

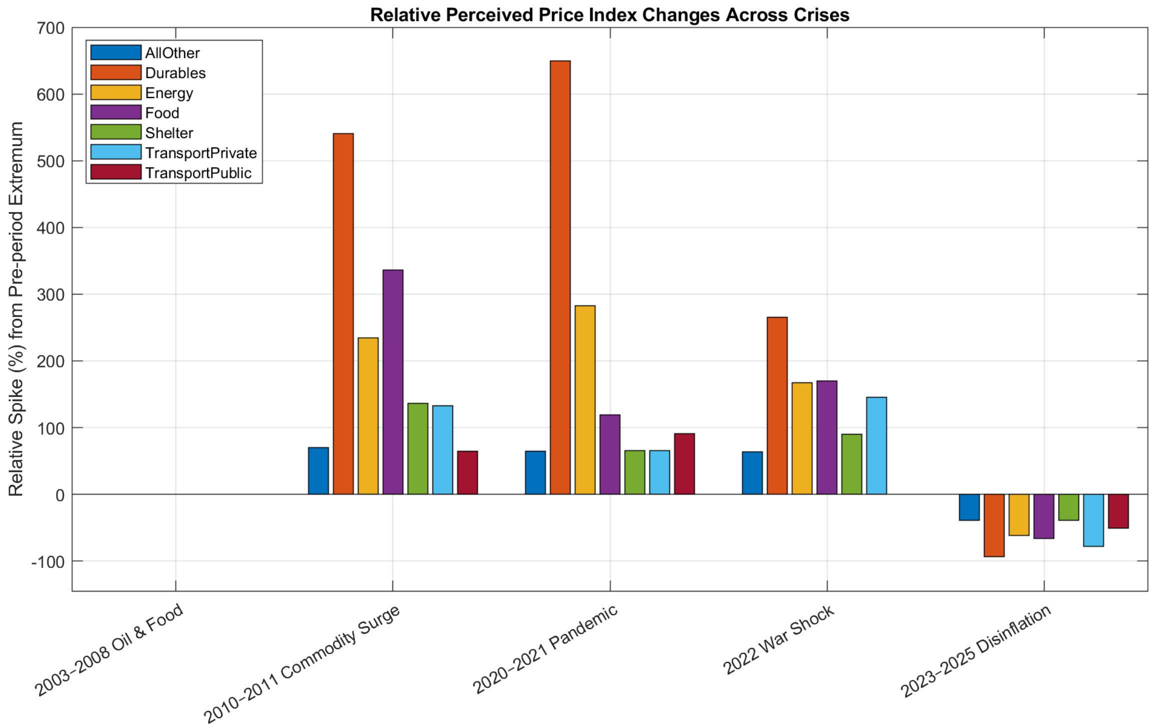

Figure 7: **Relative Index Changes Across Crises.** Grouped bars display the percent change in each category relative to its pre-crisis extremum. Positive values represent crisis-driven spikes, while the disinflation period shows negative easing.

In Figure 7, the media coverage index exhibits the largest shifts in the categories most directly linked to each respective crisis. Notably, durable consumption—though not highlighted among the primary categories above—emerges as the most intensively covered during periods of turmoil. This pattern likely reflects the heightened volatility of consumer durables under extreme economic conditions.

Obtaining Search-Based Attention Measures

To capture consumer attention through online information-seeking behavior, we construct a monthly panel of search-based measures using Google Trends.

Keyword Design and Category Mapping

Each CPI category is mapped to a curated set of 5–15 search phrases reflecting real-world language used by consumers when seeking price information. For example, the "Food" category includes terms such as *"food prices," "grocery bill," "cost of groceries,"* and *"meal prices"*. Term selection was guided by domain knowledge, exploratory search behavior, and media co-occurrence patterns to ensure relevance and behavioral plausibility.

Data Collection Using Google Trends

We use the `pytrends` Python interface to the Google Trends platform to collect normalized monthly interest scores for each term. For each term-category pair, we query U.S. search interest at monthly resolution from January 2004 to May 2025. Google Trends provides a relative index ranging from 0 to 100, where 100 represents the term's peak popularity during the query window.

To avoid rate limits and sampling inconsistencies, we restrict each query to a single-month time frame and include a delay of 1.5 seconds between requests. Errors and incomplete responses are handled with retry logic and logging.

Aggregation to Category-Level Indices

For each month, we compute the average search interest across all terms within a given CPI category. Let $GT_{i,t}$ denote the normalized search interest for term i at time t, and let c be a CPI category composed of N_c terms. The monthly category-level index is computed as:

$$\text{SearchIndex}_{c,t} = \frac{1}{N_c} \sum_{i=1}^{N_c} GT_{i,t}$$

This averaging procedure reduces the influence of idiosyncratic spikes and produces a more stable measure of collective search behavior at the category level.

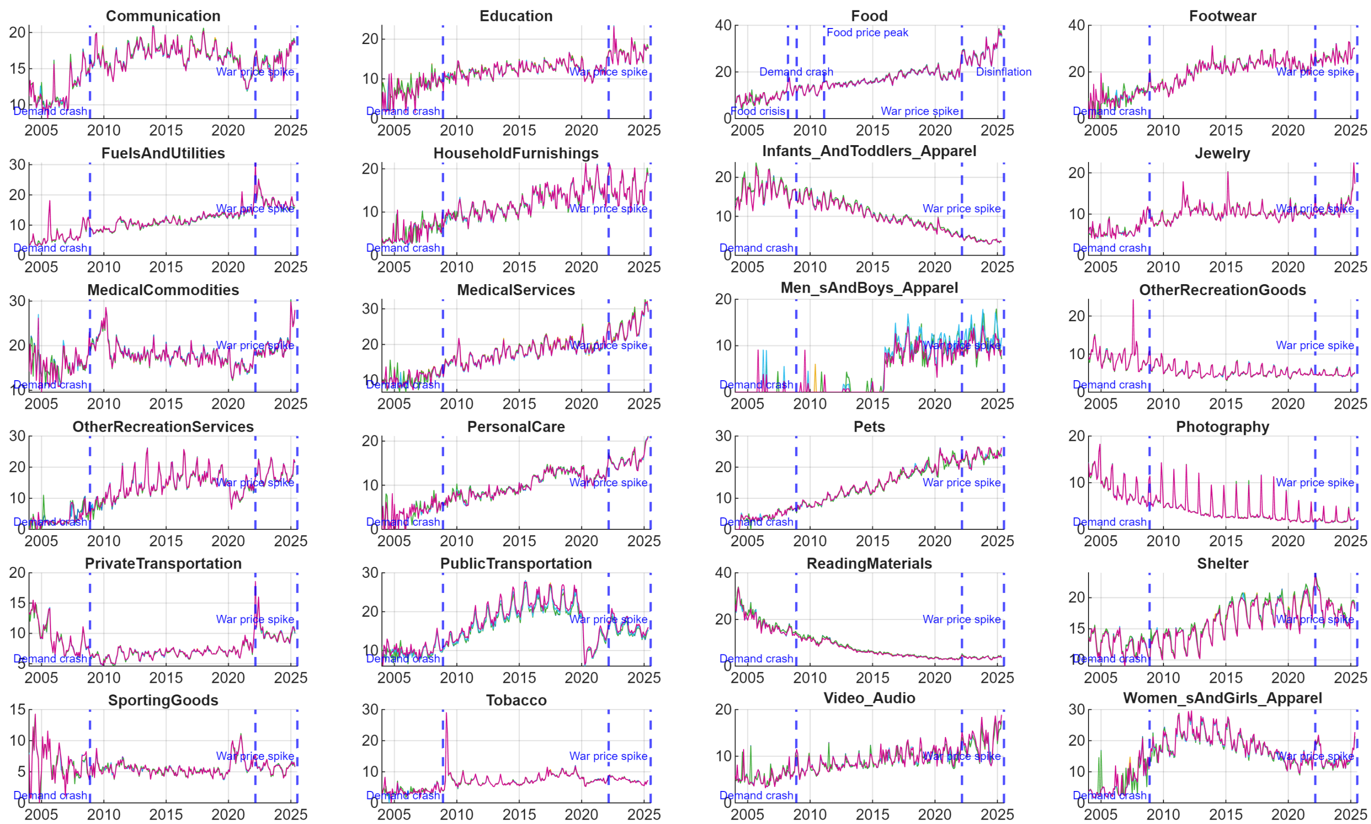

Figure 8: Results from multiple executions of the Python data processing script. Each subplot corresponds to a distinct category, with lines representing outcomes from separate runs of the code. Vertical dashed lines indicate notable events, while shaded regions mark the oil crisis period for relevant categories.