HEIDELBERG UNIVERSITY
DEPARTMENT OF ECONOMICS



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

Sector-Specific Household Inflation Uncertainty and Monetary Policy

Philip Schnattinger

Prachi Srivastava

AWI DISCUSSION PAPER SERIES NO. 772
July 2026

Sector-Specific Household Inflation Uncertainty and Monetary Policy *

Philip Schnattinger† Prachi Srivastava‡

July 16, 2026

We study the drivers of household inflation uncertainty using microdata from the Federal Reserve Bank of New Yorks Survey of Consumer Expectations from 2013 to 2025. Exploiting rounding behavior in sectoral price forecasts, we construct household-level measures of inflation uncertainty for food, gasoline, rent, medical care, and education. Combining these measures with BLS micro-price dispersion, we document a transmission channel from sectoral price dispersion to household sector-specific uncertainty and, ultimately, to aggregate inflation uncertainty. Food-price uncertainty explains the largest share of aggregate uncertainty, while gasoline uncertainty becomes increasingly important after 2020. Using identified monetary policy shocks, we show that contractionary policy lowers sector-specific inflation uncertainty and compresses price-change dispersion almost exclusively in food and gasoline. We interpret these findings through a noisy-information two-sector New Keynesian model in which price dispersion affects the precision of households price signals. With its transmission parameters estimated on these empirical responses, the model implies that optimal policy differs across moments: central banks should look through flexible food and gasoline price-level fluctuations, but lean against uncertainty originating in these sectors, while placing greater weight on price-level stabilization in stickier sectors.

Keywords: Inflation Uncertainty, Household Expectations, Monetary Policy, Food Prices
JEL Codes: D01, D81, D84, D91, E31, E52

*We thank Andrea Alati, Constantin Bürgi, Jonas Dovern, Maximilian Günnewig Mönert, Matthew Naylor, Oana Peia, Oliver Pfäuti, Philippe Andrade, Lena Dräger, Giang Nghiem, Mengdi Song, Karl Whelan, and Johannes Wohlfart for helpful comments. We also thank seminar and workshop participants at the Africa Meeting of the Econometric Society 2026, Heidelberg University, ifo Dresden workshop 2026, the Bank of England, the Bank of Slovenia, CEBRA, Workshop on Uncertainty, Economic Activity, and Policy, Warsaw, and the University of Piraeus. The views expressed are solely our own and do not reflect those of the Bank of England.

†Bank of England. philip.schnattinger@bankofengland.co.uk
‡Heidelberg University. prachi.srivastava@awi-heidelberg.de

1 Introduction

Household inflation expectations are central to modern macroeconomics, and a large body of survey-based work identifies what drives households' point forecasts.[1] Far less is known about the uncertainty surrounding those forecasts: how confident households are about price changes for the items in their consumption baskets, and how that uncertainty aggregates into overall inflation uncertainty. The gap matters. The second moment of inflation, not only its level, carries first-order welfare consequences; much of households' dislike of inflation stems from the uncertainty it creates; and randomized controlled trials document that inflation uncertainty causally shapes households' spending, saving, and portfolio decisions.[2] Yet household surveys rarely elicit item-specific uncertainties, because the probability-based questions needed to measure them are costly to field and seldom repeated.

To overcome this limitation, we adapt the round-numbers method of Binder (2017b) and apply it to households' expected price changes for individual basket items, which the New York Fed's Survey of Consumer Expectations (SCE) elicits only as point forecasts. Because uncertain respondents disproportionately report multiples of five (Krifka, 2007; Binder, 2017b), we can infer item-level subjective uncertainty from the observed responses for food, gas, medical care, education, and rent from 2013–2025. These are the categories where such uncertainty is likely to matter most: households' inflation expectations are disproportionately shaped by salient price experiences in frequently purchased categories (e.g., DAcunto et al., 2021; Coibion et al., 2023; Campos et al., 2022; Dietrich, 2024; Anesti et al., 2024). We validate the resulting uncertainty series against independent benchmarks, use them to decompose aggregate inflation uncertainty into its sector-level sources, and study how monetary policy shocks propagate through both expectations and uncertainty across consumption categories. Finally, we evaluate optimal simple monetary-policy rules that respond to item-specific inflation uncertainty.[3]

In this paper, we make four main contributions. Our first contribution is to construct and validate category-specific measures of household inflation uncertainty. Existing measures of macroeconomic uncertainty are predominantly aggregate (e.g., Bloom, 2014; Baker et al., 2016; Jurado et al., 2015). At the household level, Dominitz and Manski (1997) measure subjective uncertainty about income growth, while Ben-David et al. (2018) document heterogeneity in uncertainty about personal and macroeconomic outcomes. For inflation, Binder (2017b) constructs a proxy for subjective uncertainty using the Michigan Survey of Consumers, Reiche and Meyler (2022) constructs an inflation uncertainty index

[1] See DAcunto et al. (2024) for a recent overview and, e.g., Hoffmann et al. (2022), Andrade et al. (2023), and DAcunto et al. (2021); Weber et al. (2022) survey this evidence, documenting that household inflation expectations are driven primarily by salient and frequently-encountered prices, personal experience, demographic exposure, and the degree of attention households pay to inflation, rather than by aggregate macroeconomic conditions alone.

[2] On the welfare consequences of the second moment, see Friedman (1977) and Ball (1992); on households' dislike of inflation operating through the uncertainty it creates, Stantcheva (2024) and Binetti et al. (2024); for the causal evidence from randomized controlled trials, Georgarakos et al. (2024), Kostyshyna and Petersen (2024), and Fischer et al. (2025).

[3] The prominence of these categories is consistent with the structure of the Consumer Price Index (CPI), where food, energy, and shelter together account for over 55% of the index. According to current Bureau of Labor Statistics weights, food comprises 13.6%, energy 6.3% (with gas at 2.9%), and shelter 35.4% (including 7.5% for rent of primary residence; source: Table 1, https://www.bls.gov/news.release/cpi.t01.htm).

for the euro area, and the SCE elicits probability distributions from which uncertainty about aggregate inflation can be calculated directly (Armantier et al., 2017). Relative to this literature, we recover uncertainty about expected price changes at the level of individual consumption categories. Applying the Round-Number Reporting of Imprecision (RNRI) approach to point forecasts for gasoline, food, medical care, education, and rent, we construct individual-level measures and aggregate them into monthly category-specific uncertainty indices. Because rounding need not reflect uncertainty alone, we validate these measures against respondents' density-based inflation uncertainty, differences in rounding across categories within the same respondent, and subsequent forecast errors.[4]

Second, we quantify how uncertainty in specific consumption categories contributes to households' overall inflation uncertainty. We show that food-price uncertainty accounts for the largest share of the explained variation in aggregate inflation uncertainty, followed by rent, education, and gas. We then trace the evolution of these contributions over time using Kalman smoothing and Shapley decomposition, finding that gas-price uncertainty becomes more important after 2020. These patterns are consistent with evidence that households overweight salient price signals when forming inflation beliefs, particularly prices that are frequently observed or prominently reported, such as gasoline and grocery prices (DAcunto et al., 2021; Binder, 2018; Coibion and Gorodnichenko, 2015b; Binder, 2018). Sectoral uncertainty is also persistent, and its response to shocks builds over several months rather than resetting: uncertainty, once created, fades slowly.

Third, we provide evidence on how monetary policy transmits to household beliefs. Using the high-frequency monetary policy shocks of Jarociński and Karadi (2020) in a local-projection framework, we show that monetary policy shocks lower both inflation expectations and inflation uncertainty across categories: aggregate inflation uncertainty falls by about 0.06 standard deviations per 25 basis points, the decline builds over about five months, and the expectations response is concentrated in the most salient categories, gasoline and food. We trace this heterogeneity to a price-dispersion channel: using BLS CPI micro-data, we show that monetary policy compresses the cross-sectional dispersion of prices in flexible-price categories such as food and gas, tightening the distribution from which households draw their price signals, while dispersion in sticky-price categories changes little. Monetary policy thus anchors not only what households expect but how sure they are of it.

Fourth, we develop a structural noisy-information two-sector New Keynesian model in which a shared Kalman-filtering block over latent sectoral and demand pressures delivers a micro-founded, time-varying posterior inflation-uncertainty object. The implied posterior covariance is policy-sensitive because monetary policy compresses signal noise, so the central bank can move sectoral uncertainty through its rate decisions. We estimate the model's transmission parameters, including a cross-category salience loading in the spirit of DAcunto et al. (2021), by minimum distance on the impulse responses and dispersion moments of our third contribution; the timing and shape of the services response carry no parameter of their own and serve as over-identifying restrictions. We then ask whether and how strongly the central bank should respond to the posterior uncertainty object in an augmented Taylor rule. The model implies an inversion of the standard core-versus-

[4]Rounding may reflect reporting conventions as well as subjective imprecision (Clements, 2021; Dovern and Glas, 2025). Our validation exercises address this concern directly.

headline prescription across moments, driven by a single primitive: flexible food-and-gas prices generate little level misallocation, so the welfare-maximising rule looks through their level and targets core inflation, yet their price-change dispersion is exactly what policy can compress, so the rule leans against their uncertainty, and does the reverse for services. The lean is substantial in rate terms, about 100 basis points per standard deviation of food-and-gas uncertainty, roughly twice the response to core inflation, and ignoring uncertainty altogether costs about 25 basis points of consumption each period, the conventional headline rule about 35, several times the classic cost of business cycles.

Our normative results speak to the long-standing debate over whether monetary policy should target core or general inflation. The canonical multi-sector benchmark prescribes stabilizing a stickiness-weighted core index and looking through flexible food and energy prices (Aoki, 2001; Mankiw and Reis, 2003; Eusepi et al., 2011; Bodenstein et al., 2008), while a countervailing literature argues for attending to headline and food prices more than their distortion weights imply, because credit-constrained or inattentive households are exposed to the salient prices that core excludes (Anand et al., 2015; DAcunto et al., 2021; Bullard, 2011). We add a distinct rationale that operates on the second moment: the volatile categories that core strips out, food and gasoline, are among the dominant contributors to aggregate household inflation uncertainty, together with rent, and that uncertainty is itself policy-responsive, falling under contractionary policy most strongly in the flexible food-and-gas sector. This gives a risk-management reason to monitor general rather than core inflation, one that leaves the sticky-price logic intact, and sharpens the expectations-based case of Anesti et al. (2024) from the level of household expectations to their dispersion.

The rest of the paper proceeds as follows. Section 2 motivates how subjective inflation uncertainty shapes individuals' inflation expectations through a signal-extraction framework. Section 3 describes the dataset and our methodology for constructing sector-specific uncertainty indices. Section 4 examines the contribution of each sector to overall inflation uncertainty. Section 5 explores the response of uncertainty to monetary policy shocks. Section 6 presents the structural noisy-information New Keynesian model, estimates its transmission parameters on our empirical impulse responses, and evaluates optimal simple policy responses to sectoral inflation uncertainty under both aggregate and sector-specific rules. Section 7 concludes.

2 A Framework for Individual Inflation Expectations and Uncertainty

We motivate our empirical decomposition with a parsimonious signal-extraction framework in the spirit of Nord (2022), Fuster et al. (2022), and Maćkowiak et al. (2023). It formalises two ideas: households form inflation beliefs by combining noisy, item-specific price signals, and heterogeneity in the precision of those signals generates cross-sectional variation in both inflation expectations and inflation uncertainty. Derivations are collected in Appendix A.

Let $j \in \mathcal{J} \equiv \{\text{gas, food, rent, medical care, education}\}$ index consumption categories. Each household i holds a category-specific prior over one-year-ahead price changes, $\pi_{j,t+1} \sim \mathcal{N}(\mu_{i,j}, \sigma^2_{i,j})$, heterogeneous across households in both the perceived mean and the prior

variance. In each period t, household i observes a noisy signal about category-j price changes,

$$s_{i,j,t} \;=\; \pi_{j,t+1} \;+\; \eta_{i,j,t}, \qquad \eta_{i,j,t} \sim \mathcal{N}\Big(0,\, \sigma^2_{\eta,i,j,t}\Big). \tag{1}$$

The effective signal variance $\sigma^2_{\eta,i,j,t}$ reflects both category-specific information frictions (e.g., purchase frequency and price salience) and individual cognitive noise. A concrete and observable component is the cross-sectional dispersion of prices within the category: an item's price change equals the category component plus an idiosyncratic deviation, so when households read item price changes as signals of the category, the cross-item dispersion is part of $\sigma^2_{\eta,i,j,t}$. This ties the signal variance to realised price dispersion, a link we document empirically in Section 5.2 and embed structurally in Section 6. The two components also structure the correlations the framework predicts: cognitive noise is household-specific and common to every category the household tracks, so item-level uncertainties co-move within a household; and the dispersion component moves together across flexible-price categories, whose price distributions respond to the same aggregate conditions, so that co-movement should be strongest within the salient non-core group.

With conjugate normal priors and signals, Bayesian updating yields an item-level posterior whose precision is the sum of prior and signal precision,

$$\mathrm{Var}_i(\pi_{j,t+1} \mid s_{i,j,t})^{-1} \;=\; \sigma^{-2}_{i,j} \;+\; \sigma^{-2}_{\eta,i,j,t}, \tag{2}$$

and whose mean is the corresponding precision-weighted average of the prior $\mu_{i,j}$ and the signal $s_{i,j,t}$. Two implications follow. Noisier signals for category j raise the posterior variance of that item's price change and shift weight from the signal to the prior; and heterogeneity in priors $\{\mu_{i,j}, \sigma^2_{i,j}\}$ and in signal precisions $\{\sigma^{-2}_{\eta,i,j,t}\}$ generates cross-sectional dispersion in both posterior beliefs and posterior uncertainty at the item level, which then propagates to the basket level.

When making decisions, households care about the perceived inflation of their consumption basket as a whole. Let one-year-ahead perceived basket inflation be a fixed-weight sum of item-level price changes, $\Delta p_{i,t+1} = \sum_{j=1}^{J} \omega_j \, \Delta p_{i,j,t+1}$ with $\sum_{j=1}^{J} \omega_j = 1$, where the ω_j are perceived expenditure shares held fixed throughout. Conditional on information $\mathcal{I}_{i,t}$, household i's basket-level subjective inflation uncertainty is

$$\begin{aligned} \sigma^2_{\pi,i,t+1} &= \mathrm{Var}_i\Big(\Delta p_{i,t+1} \mid \mathcal{I}_{i,t}\Big) \\ &= \sum_{j=1}^{J} \omega_j^2 \, \sigma^2_{i,j,t+1} \;+\; 2 \sum_{1 \le j < k \le J} \omega_j \omega_k \, \mathrm{Cov}_i\Big(\Delta p_{i,j}, \Delta p_{i,k} \mid \mathcal{I}_{i,t}\Big), \end{aligned} \tag{3}$$

where $\sigma^2_{i,j,t+1} \equiv \mathrm{Var}_i(\Delta p_{i,j,t+1} \mid \mathcal{I}_{i,t})$ are the item-level perceived uncertainties of (2) and the covariance terms capture perceived co-movement across items. Equation (3) is the framework's main output: it maps item-level uncertainty and cross-item co-movement into each household's aggregate inflation uncertainty. It is the object we measure in Sections 3–4, and whose response to monetary-policy shocks we estimate in Section 5; its weights $\{\omega_j^2\}$ carry into the welfare aggregator of the model in Section 6.

Signal precision and monetary policy. The signal noise in (1) inherits the realised cross-sectional dispersion of the item price changes households sample, and that dispersion is policy-responsive: with positive trend inflation, a change in a category's inflation moves

the year-on-year dispersion of its price changes at first order, and by more in categories whose prices adjust more frequently (Appendix J.3). Tighter policy therefore compresses the distribution households sample from most in flexible-price categories: food and gasoline are more policy-sensitive than rent or education, with the ordering following from the frequency of price adjustment rather than from any assumed category-specific loading. Attention, finally, is a household-level resource: when the most visible dispersion, food and gasoline, absorbs processing capacity, signal noise rises in other categories as well, so any response of core-category uncertainty to policy operates through this cross-category spillover rather than through own-price dispersion. Section 6 embeds the framework in general equilibrium: the exogenous signal variance $\sigma^2_{\eta,i,j,t}$ becomes the model's own equilibrium price-change dispersion, and the static update (2) becomes a Kalman filter over the persistent states that drive sectoral inflation.

The framework delivers four testable predictions:

(i) Item-level uncertainties should co-move within households, most strongly among the salient non-core items, so that categories cluster into a non-core group (food, gas) and a core group (rent, medical care, education) (Section 3, Figure 5).

(ii) Basket-level uncertainty should be dominated by categories that combine large perceived weights with noisy signals, since contributions in (3) scale with $\omega_j^2\sigma_{i,j,t}^2$: salient food and large-weight rent are the natural leaders, with gas gaining whenever its signal noise spikes (Section 4).

(iii) Because signal noise inherits realised price dispersion, which evolves with the outstanding price distribution and is therefore slow-moving, uncertainty should be persistent once created (Section 3.6).

(iv) Monetary policy should move item-level uncertainty through the dispersion of the price signals households sample: contractionary policy compresses dispersion, and with it uncertainty, in flexible non-core categories where prices actually adjust; in sticky core categories dispersion itself barely responds, so any decline in core uncertainty must arrive through cross-category information spillovers rather than own-price dispersion (Sections 5 and 5.2).

3 Data and Measurement

3.1 The Survey of Consumer Expectations

We use data from the Survey of Consumer Expectations (SCE) conducted by the Federal Reserve Bank of New York (FRBNY). The SCE is a monthly rotating panel survey that collects data on consumer expectations and demographic characteristics; respondents participate for up to 12–16 consecutive months. Our analysis uses the core monthly survey from June 2013 to June 2025, covering 145 months and approximately 1,200–1,300 respondents per month.

The core survey measures inflation expectations in two formats. The first is a point forecast, in which individuals report their best guess for the one-year-ahead inflation rate (Q8). The second is a subjective probability distribution, in which respondents allocate probabilities across bins of inflation outcomes (Q9), providing a direct measure

of expected inflation uncertainty.[5] The survey asks a similar set of questions for the longer-term horizon, covering 24–36 months ahead, and separately elicits sectoral point forecasts for gas, food, rent, medical care, and education (question C4). Appendix B reproduces the exact wording of each question block, together with the response bins.

We use both point forecasts and fitted subjective density forecasts to construct our inflation expectations variables. Following Crump et al. (2022), we trim all variables at the 1st and 99th percentiles and drop observations whose reported point forecast lies outside the 1st–99th percentile range of the respondent's own fitted density forecast, since such cases indicate that the point estimate and the underlying distribution are not internally consistent. Sectoral point forecasts are trimmed at the same percentiles. From the point forecasts for these five categories, we construct our sector-specific uncertainty indices.

3.2 Density-Based Aggregate Uncertainty

To quantify subjective inflation uncertainty, we use the standard deviation of the individual's fitted subjective probability distribution over future inflation outcomes. The methodology for estimating the distribution is based on the approach in Engelberg et al. (2009), and the FRBNY SCE implements it as detailed in Armantier et al. (2017). The specific distribution fitted to each respondent's reported probabilities depends on the number of bins selected by each forecaster.

Figure 1 plots the average level of subjective inflation expectations and uncertainty over time. Before the COVID-19 pandemic, expected inflation persistently exceeded actual inflation levels. Following the pandemic, inflation expectations surged during 2021 but subsequently declined. Inflation uncertainty rose sharply, reaching 3.7% in April 2020 and peaking at 4.6% in June 2022. As of June 2025, uncertainty remains elevated at 3.2%, significantly above pre-pandemic levels.

3.3 A Sectoral Uncertainty Proxy: Round-Number Reporting

Using SCE data, we construct sector-specific uncertainty proxies. We use individual's one-year-ahead point forecasts for five consumption categories in the SCE: gas, food, medical care, education, and rent. We adopt the round-number methodology originally proposed by Binder (2017b) in the context of the Michigan Survey of Consumers. The approach builds on the insight of Krifka (2007), who introduced the concept of Round-Number Reporting of Imprecision (RNRI): individuals with greater subjective uncertainty tend to report forecasts at conventional focal values, such as multiples of five, rather than precise figures. This behavioural tendency serves as a revealed signal of an individual's confidence in their own prediction. Formally, we classify a response as a round-number (M5) forecast if it is a multiple of five (for example, -5, 0, 5, 10, 15) and as a non-round forecast otherwise (for example, 3, 7, 9). We treat respondents who report non-round values as low-uncertainty types and those who report M5 values as high-uncertainty types.

Figure 2 displays the distribution of price expectations across the five sectors, alongside the overall inflation expectations forecast. In all six panels, there are striking spikes at

[5] As noted by Zhao (2023), household forecasts tend to exhibit more dispersion and coarseness in the subjective distributions compared to professional forecasters. This reflects both higher levels of perceived uncertainty and noisier beliefs.

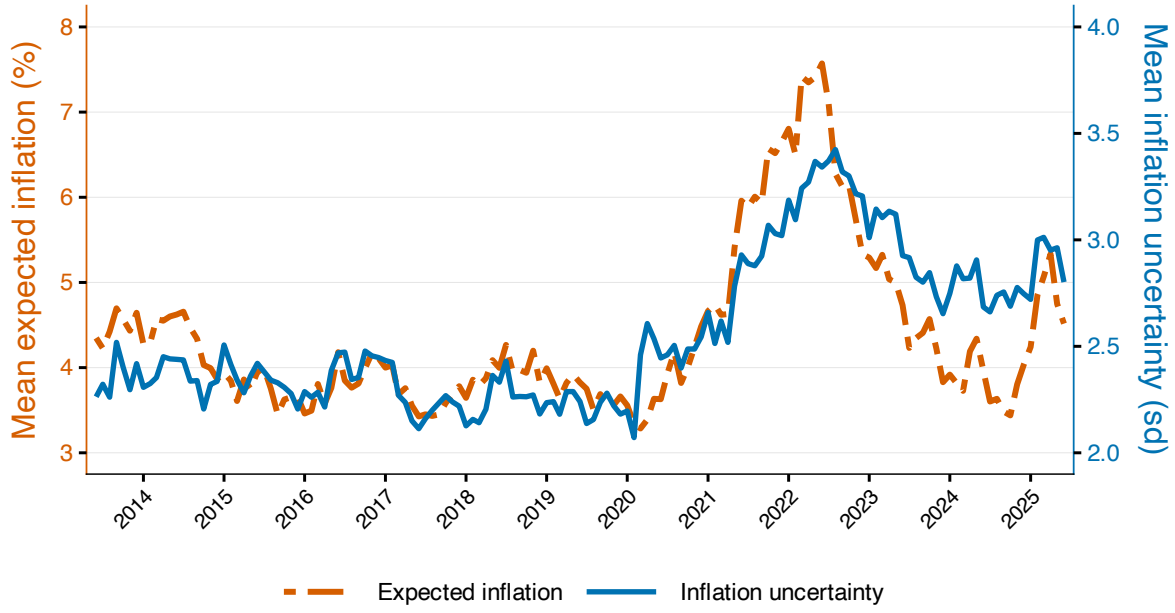


Figure 1: Inflation Expectations and Uncertainty; FRBNY SCE

Note: This figure plots the average values of mean inflation expectations and uncertainty between June 2013 and June 2025 using the FRBNY Survey of Consumer Expectations (SCE). The data is trimmed at 1-99th percentile

multiples of five and ten. We interpret this heaping as evidence of widespread rounding behaviour and, more broadly, limited reporting precision among survey respondents. The degree of heaping varies across categories in an intuitive way. Sectors for which respondents are likely to have less frequent or direct price information, such as medical care and education, exhibit sharper concentration at a few round values and heavier right tails, suggesting that greater uncertainty leads to more coarse rounding and more extreme responses. By contrast, gas and food, items purchased regularly and whose prices are readily observable, still display heaping but with relatively more mass near the centre of the distribution, consistent with respondents holding somewhat firmer beliefs about these categories.

While M5 responses are indicative of coarser beliefs, they do not always imply high uncertainty; rounding may also occur for convenience. To formally distinguish high- from low-uncertainty types, we treat the cross-sectional distribution of reported forecasts in each sector-month as a two-component mixture and estimate, by maximum likelihood, the posterior probability that a given response comes from the high-uncertainty component. Low-uncertainty types are more precise, and their reported values are modelled with a narrower support; high-uncertainty types use coarser approximations and require a wider integration range. Formally, in month t responses for sector j are generated by

$$\psi_{jt}(r) = \lambda_{jt}\,\psi^{h}_{jt}(r) + \left(1 - \lambda_{jt}\right)\psi^{l}_{jt}(r),$$

where ψ^{h}_{jt} and ψ^{l}_{jt} are the discretized normal probability mass functions of type-h and type-l respondents and $\lambda_{jt} = N^{h}_{jt}/N_{jt}$ is the share of high-uncertainty types. Given

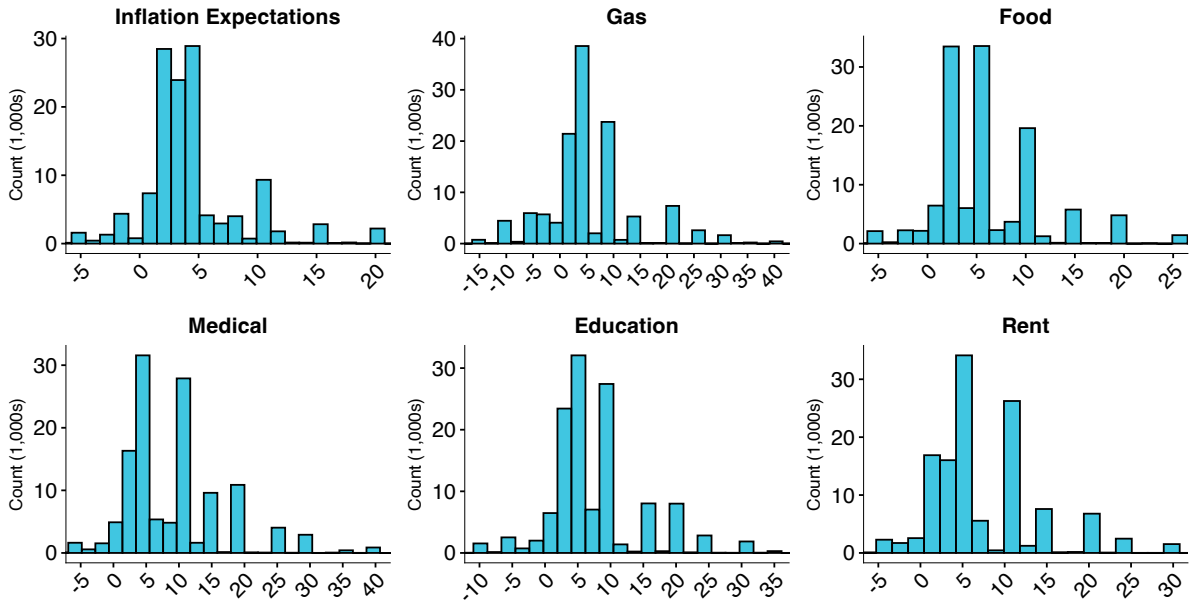


Figure 2: Inflation and Sectoral Price Expectations

Note: This figure plots FRBNY SCE respondents' forecasts of sectoral price changes. The sample covers June 2013 to June 2025 and is trimmed at the 1st and 99th percentiles.

individual responses $\{R_{jit}\}_{i=1}^{N_{jt}}$ we recover the five parameters $\{\lambda_{jt}, \mu^l_{jt}, \mu^h_{jt}, \sigma^l_{jt}, \sigma^h_{jt}\}$ by maximum likelihood; Appendix C states the likelihood and the construction of ψ^h and ψ^l.

Applying Bayes' rule to the estimated mixture gives the posterior probability that a given observation belongs to the high-uncertainty type,

$$\chi_{jit} = \frac{\lambda_{jt}\psi^h_{jt}(R_{ijt})}{\lambda_{jt}\psi^h_{jt}(R_{ijt}) + (1-\lambda_{jt})\psi^l_{jt}(R_{ijt})}, \tag{4}$$

which is our main uncertainty proxy for sector j, individual i, in month t.[6]

3.4 Validating the RNRI Measure

A limitation of any rounding-based measure is that focal-number responses are not direct observations of subjective uncertainty. Heaping at multiples of five is consistent with imprecise beliefs, but it could also reflect coarse numerical reporting or persistent survey-response habits. We therefore validate the RNRI measure using three complementary sources of evidence. First, for aggregate inflation expectations, the SCE separately elicits subjective probability distributions, allowing us to compare RNRI directly with a density-based uncertainty benchmark. Second, using the rotating-panel dimension of the SCE, we ask whether sectoral-level RNRI measures predict aggregate density-based uncertainty within the same respondent over time. Third, we examine whether high-RNRI sectors forecasts are followed by larger realized forecast errors.

[6]For details on this methodology, refer to Binder (2017a). In Appendix D we present the replication of the rounding behaviour for MCS sentiment data. We find a correlation of 98 percent between the Binder (2017b) series and our computation.

We begin with the aggregate benchmark. Following Binder (2017b), we estimate whether respondents who report rounded inflation forecasts also report wider subjective probability distributions. Let $\sigma_{it}^{\pi,\text{dens}}$ denote the standard deviation of respondent i's fitted density forecast for aggregate inflation. We estimate:

$$\sigma_{it}^{\pi,\text{dens}} = \alpha + \beta R_{it} + \varepsilon_{it}, \tag{5}$$

where R_{it} is either an indicator for reporting an aggregate inflation forecast that is a multiple of five or the RNRI-implied posterior probability χ_{it} of being a high-uncertainty type.

Table 1 reports the results. Both the multiple-of-five indicator and χ_{it} are strongly associated with density-based inflation uncertainty, but the continuous RNRI measure has substantially greater explanatory power. The M5 dummy enters positively and significantly, indicating that M5 point forecasts are associated with forecast standard deviations roughly 2.1 times as large as those of non-M5 forecasts. The proxy χ is also positive and significant, and the model including χ has a higher R^2 than those with the M5 dummy. Coefficient estimates imply that low-uncertainty types have an average forecast standard deviation of roughly 1.9%, while high-uncertainty types average around 5.8%.[7]

Table 1: Validating the relationship between RNRI and density-based uncertainty proxy

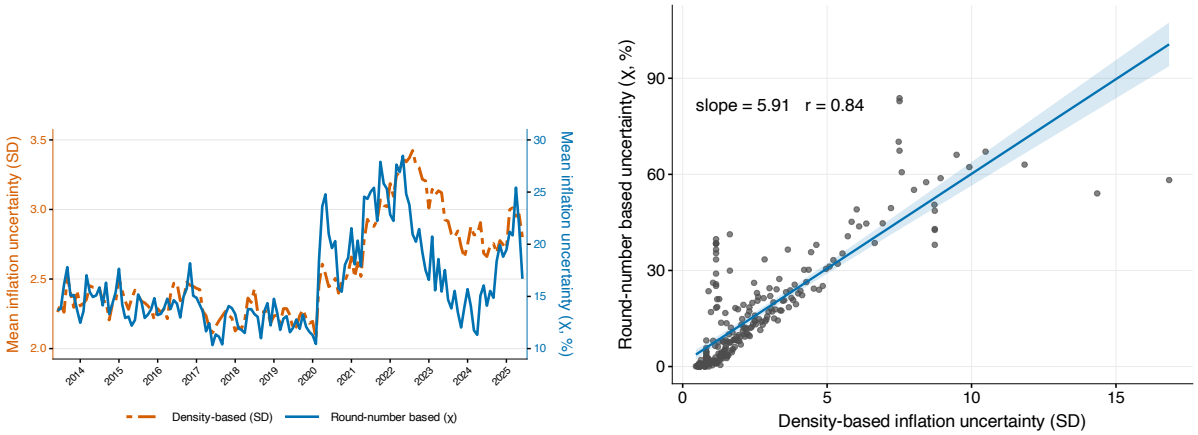
Dependent Variable:	Inflation Uncertainty	
Model:	(1)	(2)
M5 Inflation Expectations	2.091***	
	(0.0323)	
χ_{it}(RNRI)		3.851***
		(0.0493)
Constant	1.947***	1.914***
	(0.0130)	(0.0130)
Observations	127,649	127,649
R^2	0.140	0.227

Clustered (user-id) standard errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1. The dependent variable is the standard deviation of respondents' fitted density forecasts for inflation in the FRBNY SCE. M5 indicates that the point forecast is a multiple of five, and χ is the rounding-based uncertainty proxy. The data are from June 2013-June 2025.

Figures 3a and 3b provide the same comparison non-parametrically. Figure 3a shows that the monthly average RNRI index comoves closely with the average standard deviation of fitted SCE density forecasts, including during the sharp rise in uncertainty around the COVID-19 episode and the subsequent inflation surge. Figure 3b shows that this relationship also holds in the cross section: respondents with wider subjective density forecasts are systematically more likely to be classified as high-uncertainty types by the RNRI procedure.

We next test predictive validity. If χ_{jit} captures economically relevant subjective uncertainty, high-χ respondent-sector-months should have larger ex post forecast errors. Under

[7] For sectoral-level uncertainty, we estimate Equation F.6, presented in Appendix F.



(a) Time-series comparison

(b) Binned scatter plot

Figure 3: Density-Forecast Inflation Uncertainty and RNRI

Note: Panel (a) compares the monthly standard deviation implied by respondents' density forecasts with the RNRI-based uncertainty measure. Panel (b) presents a binned scatter plot of the relationship between density-based inflation uncertainty and the SCE χ measure, using 30 quantile bins of density-based inflation uncertainty. Each dot represents the average value of density-based inflation uncertainty and the average SCE χ within a bin. The fitted line is estimated using the binned observations. Sample period: June 2013 to June 2025.

full-information rational expectations with homogeneous information, the coefficient on χ_{jit} should be zero once respondent and time fixed effects are absorbed: conditional on the forecaster and survey month, RNRI type should not predict later mistakes. We estimate

$$|e_{ijt}| = \alpha_i + \delta_t + \beta_j \chi_{jit} + \varepsilon_{ijt}, \qquad e_{ijt} = \pi^{actual}_{j,t+12} - F_{ijt}, \tag{6}$$

where F_{ijt} is the reported sectoral-price forecast and $\pi^{actual}_{j,t+12}$ is the realized twelve-month-ahead price change in sector j. This test is not mechanical: χ_{jit} is inferred from the coarseness of the reported forecast, whereas the dependent variable also depends on the subsequent sectoral price realization, which is not used to construct χ_{jit}. A positive β_j therefore means that, for the same respondent in the same aggregate environment, forecasts classified as more imprecise are less accurate ex post.

Table 2 reports the estimates. Across all five categories, higher RNRI uncertainty is associated with larger realized absolute forecast errors. Since $\chi_{jit} \in [0, 1]$ is the posterior probability of being a high-uncertainty type, the coefficients compare a respondent-sector-month classified as certainly high uncertainty with one classified as certainly low uncertainty. The relationship is weakest for gasoline, the most frequently observed price category, where the high-low gap is around 2.5 percentage points, and strongest for education (5.8 pp) and medical care (6.3 pp), where prices are less transparent and less frequently encountered. The association is not a mechanical artifact of forecast levels: because high-χ respondents also report larger point forecasts, their absolute errors could be larger in a low-inflation sample regardless of subjective uncertainty, but controlling for the forecast's absolute deviation from the month-category median, or for the level itself, leaves the coefficient large and precise in every category (Appendix Table F.2).

Table 2: Predictive validity: absolute forecast errors and RNRI

Dependent Variable:			$\lvert e_{ijt} \rvert$		
Category:	Gas	Food	Medical	Education	Rent
Model:	(1)	(2)	(3)	(4)	(5)
χ_{jit}	2.501***	5.101***	6.307***	5.807***	4.286***
	(0.079)	(0.051)	(0.061)	(0.053)	(0.048)
User FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Observations	124,753	124,026	124,026	124,026	124,026
R^2	0.808	0.683	0.683	0.677	0.633

Notes: Clustered (userid) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The dependent variable is $|e_{ijt}| = |\pi^{actual}_{j,t+12} - F_{ijt}|$. Each column regresses the absolute forecast error for category j on the MLE-estimated probability of being a high-uncertainty type χ_{jit}. Sample: June 2013–June 2025.

Taken together, these exercises support interpreting χ_{jit} as a monotone proxy for relative subjective uncertainty at the sectoral level. We do not require RNRI to be a cardinal measure of posterior variance. The empirical analysis below uses it to compare relative uncertainty across households, sectors, and months.[8]

3.5 Sector-Specific Uncertainty

We next describe the sector-specific series used in the empirical analysis. We construct sector-specific uncertainty proxies, $\chi_{j,t}$, using respondents' point forecasts for each sector j and applying the RNRI method described in Section 3.3. This allows us to examine how households perceive uncertainty across different components of the consumption basket.

Figure 4 shows the evolution of average sector-specific inflation expectations and the corresponding uncertainty proxies, $\chi_{j,t}$, from June 2013 to June 2025. The figure reveals a spike in uncertainty across all sectors in March 2020, coinciding with the onset of the COVID-19 lockdowns, followed by a gradual decline. However, uncertainty remains elevated relative to pre-pandemic levels even after inflation expectations moderate, consistent with persistent effects of the pandemic and subsequent inflation episode on household price uncertainty.

Next we examine prediction (i) of the framework: do these sectoral measures capture distinct sources of uncertainty, or do they mainly move with a common household-level component, with the co-movement strongest among the salient non-core items? Figure 5 reports pooled and within-individual correlations between aggregate and sector-specific uncertainty. Panel (a) shows moderate pooled correlations between aggregate uncertainty and the sectoral measures. Panel (b) focuses on variation within respondents over time, thereby limiting the influence of persistent differences across households. The within-individual correlations with aggregate uncertainty are weaker, while correlations among the sectoral measures are stronger. This pattern suggests that households' uncertainty

[8]The online appendix reports further validation exercises; for instance, repeated SCE participation lowers both point forecasts and RNRI uncertainty, mirroring Kim and Binder (2023).

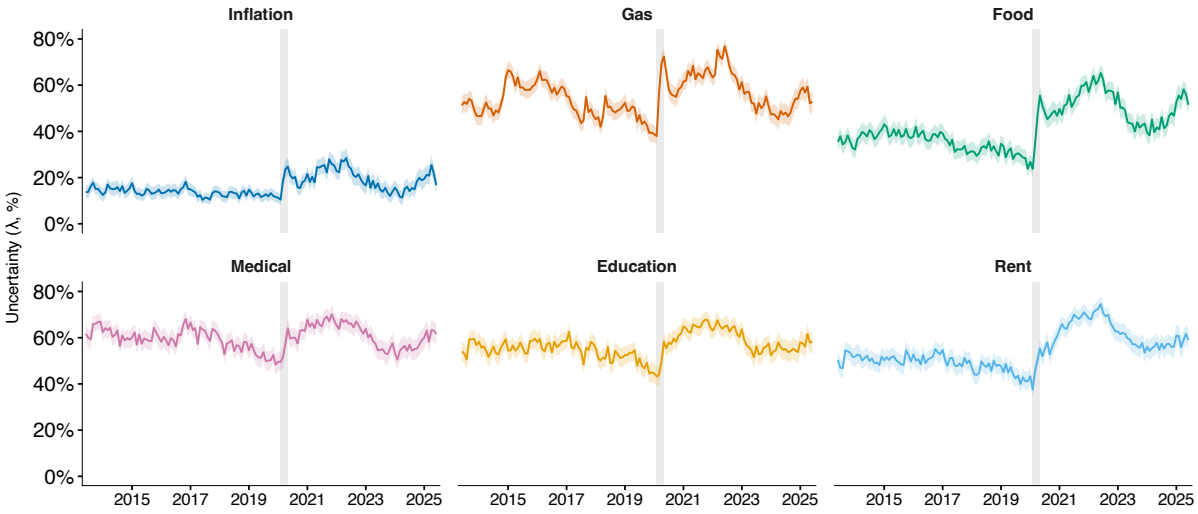
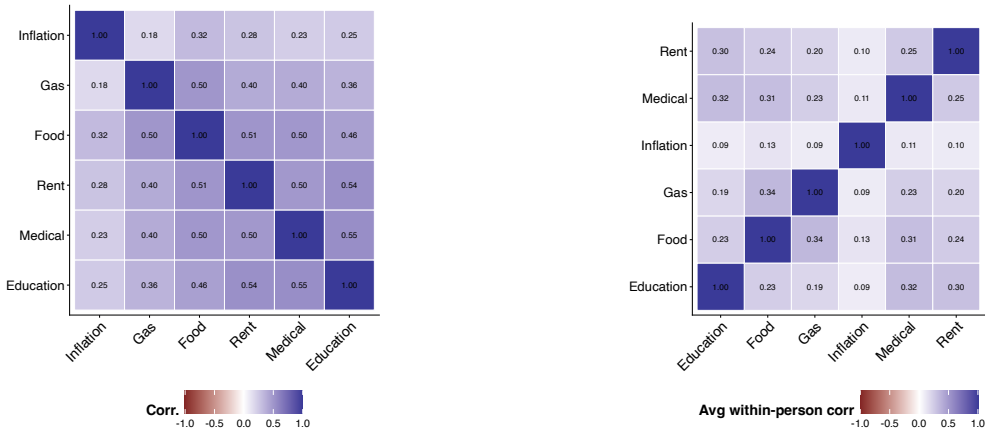


Figure 4: Sector-Specific Uncertainty Indices

Note: This figure presents the estimated monthly sector-level uncertainty indices (the share of high-uncertainty types) for five consumption categories from June 2013 to June 2025. The indices are constructed based on the authors' calculations using the Binder (2017b) method. The shaded region is the NBER COVID period from Feb 2020 to April 2020. Appendix E also presents the estimated mean and variance for each fitted distribution.

about specific consumption categories moves jointly over time, while still leaving meaningful sectoral variation for the analysis below.



(a) Pooled correlations

(b) Within-individual correlations

Figure 5: Correlations Between Aggregate and Sector-Specific Inflation Uncertainty

Note: Panel (a) reports pooled pairwise correlations between aggregate inflation uncertainty (χ) and sector-specific uncertainty measures. Panel (b) reports average within-individual pairwise correlations computed across repeated survey observations and then averaged across respondents. Sample period: June 2013 to June 2025.

3.6 Persistence of Uncertainty

We next examine prediction (iii) of the framework: whether sectoral uncertainty is persistent at the household level

$$y_{i,t}^k = \rho_k y_{i,t-1}^k + \mathbf{X}_{i,t}'\beta_k + \tau_{it} + \mu_t + \varepsilon_{i,t}^k \tag{7}$$

where $y_{i,t}^k$ denotes, alternatively, aggregate inflation uncertainty, the point forecast of inflation, or sector-specific uncertainty χ^k. The vector $\mathbf{X}_{i,t}$ includes demographic controls, τ_{it} is the tenure-id fixed effects, and μ_t denotes year-month fixed effects. Standard errors are clustered at the respondent and month level. For each outcome $y_{i,t}^k$, we regress the current value on its one-month lag, controlling for respondent characteristics. The coefficient on the lagged dependent variable, ρ_k, captures the degree of persistence net of aggregate time variation and observable household heterogeneity.

Table 3: Persistence of Aggregate and Sector-Specific Inflation Uncertainty

	Aggregate Beliefs		Sector-Specific Uncertainty				
Dependent variable:	$\sigma_{i,t}^{\pi}$	$E_{i,t}[\pi]$	$\chi_{i,t}^{\text{gas}}$	$\chi_{i,t}^{\text{food}}$	$\chi_{i,t}^{\text{medical}}$	$\chi_{i,t}^{\text{education}}$	$\chi_{i,t}^{\text{rent}}$
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Lagged dependent variable	0.693***	0.560***	0.463***	0.504***	0.483***	0.473***	0.479***
	(0.007)	(0.008)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Tenure fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	106,648	106,648	106,648	106,648	106,648	106,648	106,648
R^2	0.575	0.407	0.269	0.311	0.255	0.251	0.279

Notes: The table reports AR(1) persistence estimates for aggregate inflation uncertainty, point inflation expectations, and sector-specific uncertainty measures. Each column estimates equation 7, where $y_{i,t}^k$ denotes the dependent variable listed in the column heading. The vector $\mathbf{X}_{i,t}$ includes controls for marital status, gender, Hispanic origin, spouse, children over 18, other relatives, race, age, education, children under 18, numeracy, homeownership, state, and survey tenure. All specifications include year-month fixed effects. Standard errors are two-way clustered by respondent and month and are reported in parentheses. Significance levels are $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

The estimates indicate substantial persistence in both aggregate and sector-specific beliefs. Aggregate inflation uncertainty is the most persistent measure, with an autoregressive coefficient of 0.693, followed by point inflation expectations, with a coefficient of 0.560. The sector-specific uncertainty measures are moderately persistent, with coefficients ranging from 0.463 for gasoline to 0.504 for food. Thus, sector-level uncertainty is not merely transitory: households reporting greater uncertainty about a category in one month tend to remain relatively uncertain about that category in the following month.

3.7 Within-Household Changes in Sectoral and Aggregate Uncertainty

Having established that the sectoral uncertainty measures contain persistent and category-specific variation, we next examine whether changes in a household's uncertainty about individual consumption categories are associated with changes in its density-based aggre-

gate inflation uncertainty. We estimate the first-difference specification

$$\Delta\sigma^{\pi}_{i,t} = \beta_{\text{gas}}\Delta\chi^{\text{gas}}_{i,t} + \beta_{\text{food}}\Delta\chi^{\text{food}}_{i,t} + \beta_{\text{medical}}\Delta\chi^{\text{medical}}_{i,t} + \beta_{\text{education}}\Delta\chi^{\text{education}}_{i,t} + \beta_{\text{rent}}\Delta\chi^{\text{rent}}_{i,t} + \mathbf{X}'_{i,t}\gamma + \mu_t + \varepsilon_{i,t}. \quad (8)$$

Here, $\Delta\sigma^{\pi}_{i,t}$ denotes the change in respondent i's density-based aggregate inflation uncertainty, while $\Delta\chi^{j}_{i,t}$ denotes the change in the respondent's RNRI-based uncertainty about category j. The vector $\mathbf{X}_{i,t}$ are the demographic controls, and μ_t denotes month fixed effects.

Table 4: Changes in Sectoral Uncertainty and Changes in Aggregate Inflation Uncertainty

Dependent variable:	$\Delta\sigma^{\pi}_{i,t}$	
	(1)	(2)
$\Delta\chi^{\text{gas}}_{i,t}$	0.079***	0.073***
	(0.015)	(0.015)
$\Delta\chi^{\text{food}}_{i,t}$	0.133***	0.131***
	(0.017)	(0.017)
$\Delta\chi^{\text{medical}}_{i,t}$	0.034**	0.033**
	(0.015)	(0.015)
$\Delta\chi^{\text{education}}_{i,t}$	0.056***	0.055***
	(0.015)	(0.015)
$\Delta\chi^{\text{rent}}_{i,t}$	0.067***	0.064***
	(0.015)	(0.015)
Demographic controls	Yes	Yes
Tenure fixed effects	Yes	Yes
Year-month fixed effects	No	Yes
Observations	106,648	106,648
R^2	0.008	0.010

Notes: The table reports first-difference regressions of changes in aggregate inflation uncertainty on changes in sector-specific RNRI uncertainty. The dependent variable, $\Delta\sigma^{\pi}_{i,t}$, is the respondent-level change in aggregate inflation uncertainty from month $t-1$ to month t. The main regressors are respondent-level changes in sectoral uncertainty, $\Delta\chi^{j}_{i,t}$, for gasoline, food, medical care, education, and rent. Column (1) includes demographic controls, state fixed effects, and survey-tenure fixed effects; column (2) additionally includes year-month fixed effects. Standard errors are reported in parentheses. Significance levels are $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table 4 shows that increases in each of the five sector-specific uncertainty measures are associated with increases in density-based aggregate inflation uncertainty. The association is strongest for food: moving from the low-uncertainty to the high-uncertainty RNRI type is associated with a 0.131 increase in aggregate uncertainty in the specification with year-month fixed effects. The coefficients on gasoline, rent, education, and medical-care uncertainty are also positive and statistically significant. Their stability after the inclusion of year-month fixed effects indicates that the associations are not explained solely by aggregate developments affecting all households in a given month. Instead, changes in a household's uncertainty about particular consumption categories are systematically associated with changes in its overall inflation uncertainty.

As a further plausibility check, Appendix H documents that demographic groups with greater historical exposure to price volatility, older cohorts in particular, report systematically higher RNRI-based uncertainty, consistent with experience-based learning

(Malmendier and Nagel, 2016).

4 Accounting for Aggregate Inflation Uncertainty

This section decomposes overall density-implied inflation uncertainty into its sectoral components. We first estimate reduced-form regressions that relate household-level aggregate uncertainty to RNRI uncertainty in five consumption categories, allowing the associations to differ in the high-uncertainty period. We then let these associations vary month by month using a Kalman-smoother framework. Finally, because uncertainty across sectors co-moves, we use a Shapley decomposition to assign the explained variation in aggregate uncertainty across sectors.

4.1 Sector-Specific Contributions to Aggregate Inflation Uncertainty

The framework in Section 2 suggests that uncertainty about individual price categories contributes to households aggregate inflation uncertainty, with the largest contributions from categories combining large weights with noisy signals (prediction (ii)). We examine this empirical relationship by relating household-level aggregate inflation uncertainty to RNRI-based uncertainty measures for gasoline, food, medical care, education, and rent. We estimate the following baseline specification:

$$y_{i,t} = \alpha + \beta_1 \chi_{i,t}^{gas} + \beta_2 \chi_{i,t}^{food} + \beta_3 \chi_{i,t}^{med} + \beta_4 \chi_{i,t}^{edu} + \beta_5 \chi_{i,t}^{rent} + \gamma D_t^H + \sum_j \delta_j \left(\chi_{i,t}^{j} \times D_t^H \right) + \Lambda \mathbf{x}_{it} + \varepsilon_{it}, \quad (9)$$

Here $y_{i,t}$ is household-level aggregate inflation uncertainty, measured either from the fitted density forecast or from the aggregate RNRI proxy. The variables $\chi_{i,t}^{j}$ are sector-specific RNRI uncertainty measures for $j \in \{gas, food, medical, education, rent\}$. The high-uncertainty indicator D_t^H equals one from March 2020 onward and zero otherwise, and the interaction terms allow the sectoral associations to differ during the COVID period. All uncertainty variables are standardized, so the coefficients are interpreted in standard-deviation units.[9] The vector $\mathbf{x}_{it}$ includes demographic, housing, tenure, income, education, numeracy, state, and household-composition controls; standard errors are clustered at the respondent level.

Table 5 reports the results. In the baseline period, food uncertainty is the strongest predictor of aggregate inflation uncertainty across both dependent variables. A one-standard-deviation increase in food-price uncertainty is associated with a 0.12 standard deviation increase in density-based inflation uncertainty and a 0.18 standard deviation increase in RNRI-based aggregate uncertainty. Rent is the second-largest predictor, followed by education and gasoline. Medical-care uncertainty is not statistically related to

[9] Because $\chi_{i,t}^{j}$ is a posterior probability of belonging to the high-uncertainty (M5) type, it lies in $[0, 1]$ and is typically bimodal, with mass concentrated near 0 and near 1. Much of its standard deviation therefore reflects the gap between M5 and non-M5 respondents rather than a smooth gradient, so a one-standard-deviation change should be interpreted as approximately a shift across the M5 boundary rather than a marginal movement along a continuous scale.

density-based uncertainty, but is positive and significant when aggregate uncertainty is measured using the RNRI proxy.

The interaction terms show that these associations change during the high-uncertainty COVID period. The coefficient on D_t^H is positive and significant in both columns, indicating a higher average level of aggregate uncertainty after March 2020. Sectoral associations also shift: gasoline and medical-care uncertainty become more strongly related to aggregate uncertainty, with the clearest effects in the RNRI-based specification. By contrast, the marginal association of food uncertainty weakens in the density-based specification. This does not mean that food-price uncertainty became unimportant. Rather, COVID initially made food prices and food availability a broad common concern, as restaurant closures and stay-at-home orders sharply shifted spending toward food at home before food-away-from-home spending returned to pre-pandemic levels by April 2021; later increases in food expenditures reflected broad food-price inflation and household substitution toward more economical products.[10] In this setting, the aggregate high-uncertainty-period shift may partly absorb food uncertainty, leaving less cross-household residual variation for food to explain. The association of rent uncertainty also weakens in the RNRI-based specification.

4.2 Time-Varying Contributions to Explained Variation

The baseline specification summarizes average associations over broad periods. We next ask whether the sectoral sources of aggregate uncertainty change month by month. The price categories that households use as signals of future inflation may shift with the macroeconomic environment. We estimate time-varying sectoral coefficients in two steps, in the spirit of Fama and MacBeth (1973). In the first step, we run the following regression separately in each monthly cross-section:

$$\tilde{y}_{i,t} = \sum_{j=1}^{5} \gamma_j^t \tilde{\chi}_{ijt} + \delta_t' \mathbf{d}_{it} + \epsilon_{it} \tag{10}$$

where $\tilde{y}_{i,t} = \log(\sigma_{i,t} + 1)$, standardized to mean zero and unit variance, is the log-transformed measure of expected aggregate inflation uncertainty for individual i in month t;[11] $\tilde{\chi}_{ijt}$ is the standardized uncertainty proxy for sector j, and $\mathbf{d}_{it}$ is the vector of demographic controls (age, income, tenure, state, gender, education, marital status, children, numeracy, and race). Thus, γ_j^t measures the month-t conditional association between uncertainty about sector j and aggregate inflation uncertainty.

Monthly OLS estimates are noisy, especially in months with smaller cross-sections. In the second step, we therefore treat the vector of OLS estimates, $\hat{\boldsymbol{\gamma}}_t$, as a noisy observation of a latent smoothly evolving coefficient vector, $\boldsymbol{\gamma}_t$, and estimate the following state-space system.

[10] See USDA Economic Research Service, "Food and Consumers," https://www.ers.usda.gov/covid-19/food-and-consumers.

[11] We log-transform $\sigma_{i,t}$ prior to standardization to address the right-skewness of the uncertainty measure. The transformation $\log(\sigma_{i,t} + 1)$ also makes the cross-sectional regression residuals and the Kalman measurement noise closer to Gaussian.

Table 5: Category-Specific Uncertainty and Aggregate Inflation Uncertainty

	Dependent Variable	
Regressor	**Density-Based Inflation Uncertainty** $\widetilde{SD}_{it}^{\pi}$	**Rounding-Based Inflation Uncertainty** $\widetilde{\chi}_{it}^{\pi}$
Category-specific uncertainty		
Gas uncertainty	0.020***	0.030***
	(0.005)	(0.004)
Food uncertainty	0.120***	0.180***
	(0.006)	(0.006)
Medical uncertainty	-0.005	0.020***
	(0.005)	(0.004)
Education uncertainty	0.030***	0.060***
	(0.005)	(0.004)
Rent uncertainty	0.060***	0.100***
	(0.005)	(0.005)
Interaction with high-uncertainty period		
Gas uncertainty $\times D_t^H$	0.010	0.030***
	(0.009)	(0.007)
Food uncertainty $\times D_t^H$	-0.020**	-0.002
	(0.009)	(0.009)
Medical uncertainty $\times D_t^H$	0.020**	0.030***
	(0.009)	(0.007)
Education uncertainty $\times D_t^H$	0.004	-0.010
	(0.009)	(0.007)
Rent uncertainty $\times D_t^H$	-0.003	-0.030***
	(0.009)	(0.008)
High-uncertainty period, D_t^H	0.190***	0.110***
	(0.010)	(0.010)
Demographic fixed effects	Yes	Yes
Tenure fixed effects	Yes	Yes
Respondent-clustered SEs	Yes	Yes
Observations	123,901	123,901
R^2	0.190	0.213

Notes: This table reports regressions of aggregate inflation uncertainty on category-specific uncertainty measures and their interactions with a high-uncertainty-period indicator. All variables are standardized. The sample covers SCE data from June 2013 to June 2025. Column (1) uses the standardized density-based measure, $\widetilde{SD}_{it}^{\pi}$, as the dependent variable. Column (2) uses the standardized rounding-based measure, $\widetilde{\chi}_{it}^{\pi}$. The variable D_t^H is an indicator for the high-uncertainty period. All specifications include fixed effects for marital status, gender, Hispanic origin, household composition, race, age, education, household income, numeracy, state, homeownership, and tenure status. Standard errors clustered at the respondent level are reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

The state equation is a random walk,

$$\boldsymbol{\gamma}_t = \boldsymbol{\gamma}_{t-1} + \mathbf{w}_t, \qquad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0},\, \mathbf{W}), \tag{11}$$

and the observation equation treats the monthly OLS estimate as the state plus measurement error,

$$\hat{\boldsymbol{\gamma}}_t = \boldsymbol{\gamma}_t + \mathbf{v}_t, \qquad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \hat{V}_t), \tag{12}$$

with $\hat{V}_t$ set to the first-step covariance matrix, so months with less precise cross-sectional estimates are automatically down-weighted. We estimate the state innovation covariance $\mathbf{W}$ by maximum likelihood, using the unrestricted parameterization as our main speci-

fication so the sectoral contributions can co-move over time.[12] We report the smoothed estimate $\boldsymbol{\gamma}_{t|T}$, which conditions on the full sample. Figure 6 reports both the narrow Kalman posterior bands, shown as dotted lines, and wider bootstrap bands, shown as shaded ribbons; the former are conditional on estimated hyperparameters $(\mathbf{W}, \hat{V}_t)$ and therefore understate full parameter uncertainty.

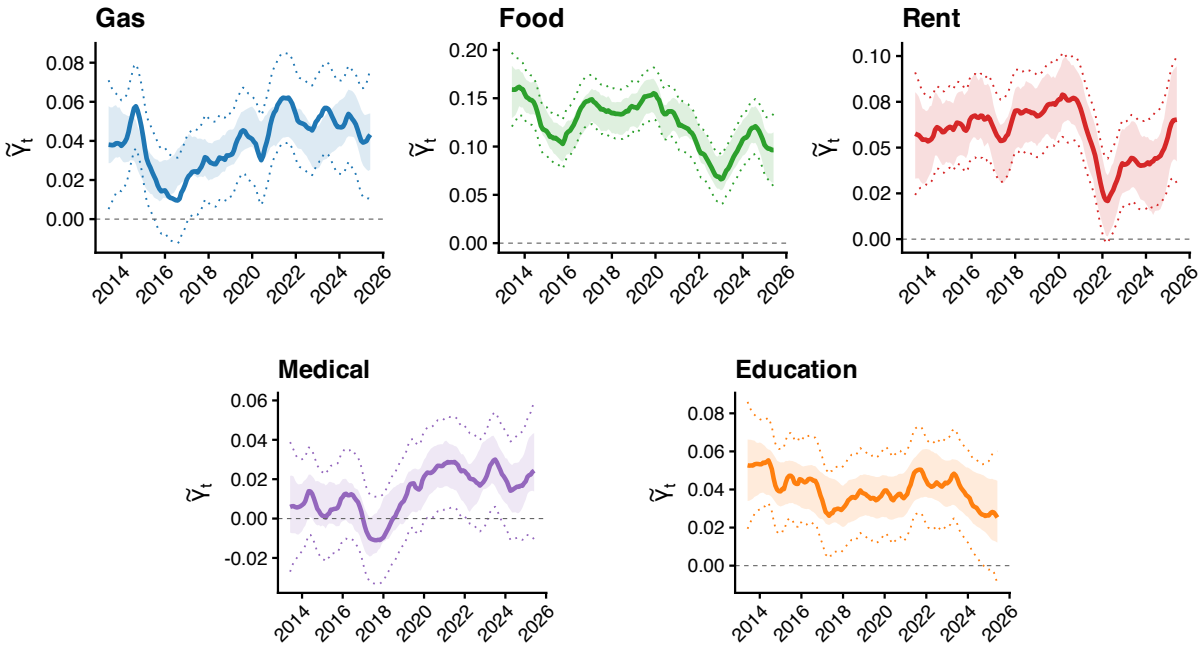


Figure 6: Time-Varying Sectoral Contributions to Aggregate Inflation Uncertainty

Note: Each panel reports the Kalman-smoothed coefficient $\gamma_{j,t|T}$ from monthly cross-sectional regressions of standardized aggregate inflation uncertainty on standardized sector-specific uncertainty, controlling for demographics. The dotted lines show conditional Kalman posterior bands; the shaded ribbons show bootstrap bands. Sample: June 2013–June 2025.

Figure 6 shows that the sectoral mapping into aggregate uncertainty is not stable over time. Food-price uncertainty has the largest and most persistent positive association before the pandemic, but its coefficient declines after 2021 and only partly recovers later in the sample. Gasoline uncertainty starts from a smaller role, rises after 2019, and remains elevated during the period of large energy-price movements. Rent uncertainty is positive throughout, falls around 2022, and then rebounds. Education uncertainty is positive but gradually weaker toward the end of the sample, while medical-care uncertainty becomes more relevant around the pandemic period, although with wider uncertainty bands. These patterns follows in the direction of pooled results: food remains central, but the marginal association of aggregate uncertainty shifts toward gasoline and medical care in the high-uncertainty period.

4.2.1 Shapley Decomposition

The Kalman coefficients describe marginal associations, but they do not by themselves allocate explained variation across sectors. This matters because sectoral uncertainty measures move together: for example, a month in which food uncertainty is high may also

[12] A likelihood ratio test against the diagonal restriction $\mathbf{W} = \text{diag}(w_1, \ldots, w_5)$ rejects independent drift.

be a month in which gasoline or rent uncertainty is high. A coefficient path can therefore show which variables are conditionally important, while still leaving open how much of the total explanatory power should be assigned to each correlated sector. To address this attribution problem, we compute a Shapley decomposition of monthly explained variation (Shapley, 1953; Shorrocks et al., 2013).

For each month t, we first residualize aggregate uncertainty and the five sectoral uncertainty proxies on the demographic controls $\mathbf{d}_{it}$. We then estimate all possible subset regressions using the residualized variables. The Shapley value for sector j is the average marginal increase in R^2 from adding that sector across all possible orderings of the five sectoral regressors:

$$\phi_{j,t} = \sum_{\mathcal{S} \subseteq \mathcal{J} \setminus \{j\}} \frac{|\mathcal{S}|!\,(K - |\mathcal{S}| - 1)!}{K!} \left[R^2(\mathcal{S} \cup \{j\},\, t) - R^2(\mathcal{S},\, t) \right], \tag{13}$$

where $\mathcal{J}$ is the set of five sectors, $K = 5$, and $R^2(\mathcal{S}, t)$ is the in-sample R^2 from the regression containing only the sectors in subset $\mathcal{S}$. Because adding a regressor cannot reduce in-sample R^2, each marginal contribution is non-negative, and the Shapley values sum to the total sectoral R^2 in month t. We normalize $\phi_{j,t}$ by $\sum_j \phi_{j,t}$ to obtain shares that sum to one.

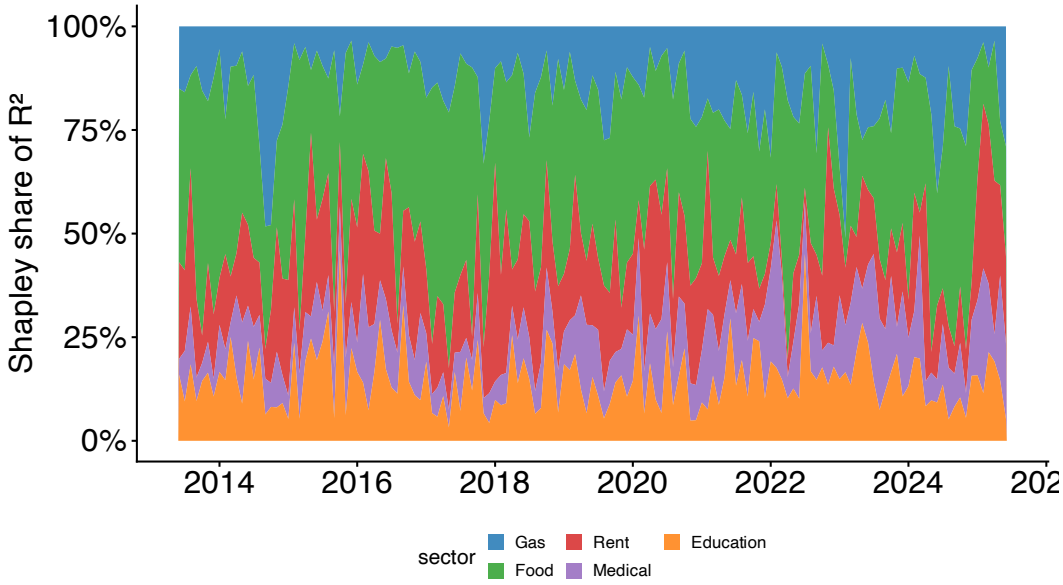


Figure 7: Shapley Decomposition of Explained Variation in Aggregate Inflation Uncertainty

Note: The figure reports monthly Shapley shares of the sectoral R^2 from regressions of aggregate inflation uncertainty on the five sector-specific uncertainty measures. Both the outcome and sectoral predictors are residualized on demographic controls within each month before computing the decomposition. Shares sum to one in each month. Sample: June 2013–June 2025.

Figure 7 confirms that food uncertainty accounts for the largest share of explained variation through much of the sample, especially before the high-inflation period, with gasoline receiving a much larger share after 2020 and rent gaining importance in several later months. The decomposition adds what the smoother cannot: an allocation of the total explained variation across correlated sectoral signals.

5 Monetary Policy Shocks: Effects on Expectations and Uncertainty

The previous section showed that households build their inflation uncertainty from uncertainty about specific price categories, especially salient categories such as food, gasoline, and rent. We now ask whether monetary policy moves these belief distributions, prediction (iv) of the framework. If policy moves the uncertainty households attach to future prices, and not just their expected level, it commands a second margin of belief stabilization. If the response is concentrated in particular sectors, then sectoral uncertainty provides a channel through which policy shapes aggregate inflation beliefs.

To identify these effects, we use high-frequency monetary policy shocks from Jarociński and Karadi (2020) and estimate local projections at the household level. These shocks provide plausibly exogenous variation in the monetary policy stance, allowing us to trace the response of aggregate inflation expectations, sector-specific expectations, and inflation uncertainty. We estimate the local projections on the pooled sample of tightening and easing shocks, so the responses below are symmetric average effects of a monetary policy shock. Appendix G shows that the results are robust to the alternative monetary policy surprise series of Bauer and Swanson (2023).

5.1 Responses of Household Expectations and Uncertainty

We estimate the following local projection over June 2013 to September 2024 (shock months through January 2024, with outcome horizons extending up to eight months beyond) using the monthly shock series of Jarociński and Karadi (2020)[13], scaling responses to a 25 basis point shock. The dependent variable is the one-year-ahead forecast, and we control for the lagged individual outcome and individual fixed effects. This specification captures the response of the level of expectations or uncertainty at each horizon, conditional on pre-shock beliefs. We estimate levels rather than changes because individual-level changes in uncertainty are noisy and can obscure persistent shifts in subjective uncertainty. To prevent the March 2020 easing surprise and the concurrent pandemic and oil-price-war shocks from contaminating the responses, the specification includes outcome-date (calendar-month) fixed effects for the acute COVID window, March through June 2020; these absorb the pandemic variation from the horizon-h outcomes without dropping any observation, so the rotating panel stays intact (Appendix I).

$$
\begin{aligned}
y_{i,t+h} = \beta_h \text{shock}_t + \sum_{l=1}^{p} \gamma_{h,l} \text{shock}_{t-l} + \delta_h' \mathbf{X}_t + \sum_{\ell=1}^{p} \phi_{h,\ell}' \mathbf{X}_{t-\ell} + \theta_h' \mathbf{Z}_t + \sum_{\ell=1}^{p} \kappa_{h,\ell}' \mathbf{Z}_{t-\ell} \\
+ \eta_h S_t + \sum_{\ell=1}^{p} \psi_{h,\ell} S_{t-\ell} + \rho_h \, y_{i,t-1} + \alpha_i \\
+ \sum_{m \in \mathcal{C}} \mu_{h,m} \mathbb{1}[t+h=m] + \lambda_{h,\text{moy}(t+h)} + \varepsilon_{i,t,h}, \qquad h = 0, \ldots, H
\end{aligned} \tag{14}
$$

where: $y_{i,t+h}$ is the individual-level outcome at horizon h, where $h \in \{0, 1, \ldots, 8\}$. For the uncertainty specifications, this is the future level of inflation uncertainty or sector-specific

[13] Appendix G plots the "poor man's" monetary policy shock series.

uncertainty. α_i are individual fixed effects, which absorb time-invariant differences across respondents. The indicators $\mathbb{1}[t+h=m]$ for $m \in \mathcal{C} = \{\text{March}, \ldots, \text{June 2020}\}$ are the outcome-date COVID fixed effects described above, and $\lambda_{h,\text{moy}(t+h)}$ are month-of-year fixed effects at the outcome date, which absorb category seasonality. The effect of the monetary policy shock is given by β_h. $\mathbf{X}_t$ is a vector of monthly macroeconomic controls, such as CPI inflation year-on-year, industrial production growth year-on-year, and the federal funds rate. The terms $\mathbf{X}_{t-\ell}$ denote their monthly lags. $\mathbf{Z}_t$ contains monthly mean inflation expectation controls with lags. We include these when estimating the response of inflation uncertainty, sectoral price expectations, and sectoral uncertainty. S_t is the sector-specific monthly mean expectation corresponding to the outcome (its lags $S_{t-\ell}$ enter likewise). $y_{i,t-1}$ is the respondent's pre-shock lagged outcome. Including this term means that β_h measures the effect of the shock on future uncertainty conditional on the respondent's initial level of uncertainty. Standard errors are clustered by survey year-month, which accounts for the within-period correlation induced by the aggregate shock.

Figure 8 reports the impulse responses from specification (14), scaled to a 25-basis-point monetary policy shock. Overall inflation expectations display a modest, short-lived decline. The sharpest adjustment is in the salient flexible-price categories: expenditure-weighted food-and-gas expectations fall by about 0.4 percentage points in the impact quarter and recover with a mild overshoot by the third quarter; gasoline is the largest but least precisely estimated component. Medical, education, and rent expectations respond later and more weakly: weighted services expectations reach about -0.3 percentage points in the second quarter and fade thereafter. The pattern is consistent with households updating beliefs most strongly about frequently observed, flexibly priced categories: households learn about inflation largely from the prices they see rather than from policy announcements, to which they are inattentive (Coibion et al., 2022; Lamla and Vinogradov, 2019; Cavallo et al., 2017).

Figure 9 reports the responses of aggregate and sector-specific inflation uncertainty to a 25-basis-point monetary policy shock. All uncertainty measures are standardized, so the coefficients are measured in standard-deviation units. Aggregate inflation uncertainty declines gradually after the shock, reaching its lowest point around five months later. The largest declines occur for gasoline, food, medical care, and rent, with troughs appearing around four to six months after the shock, while education uncertainty responds weakly. The estimates are less precise at longer horizons, but the pattern suggests the shock moves the second moment of household beliefs, not only the first. Two further checks harden these estimates: dropping the March 2020 observation from the shock series entirely leaves the responses intact or slightly deeper, and placebo projections of pre-shock outcomes on the shock show no pre-trends (Appendix I, Table I.7).

5.2 How Monetary Policy Transmits to Inflation Uncertainty

The survey responses show that monetary policy reduces uncertainty in several categories. We next ask which price-based channel can account for this response. The signal-extraction logic of Section 2 points to realized price-change dispersion. Households observe item-level price changes as noisy signals of future inflation; when the cross-sectional dispersion of those price changes narrows, the signal becomes less noisy and posterior uncertainty falls. This channel should be strongest in flexible, frequently observed cate-

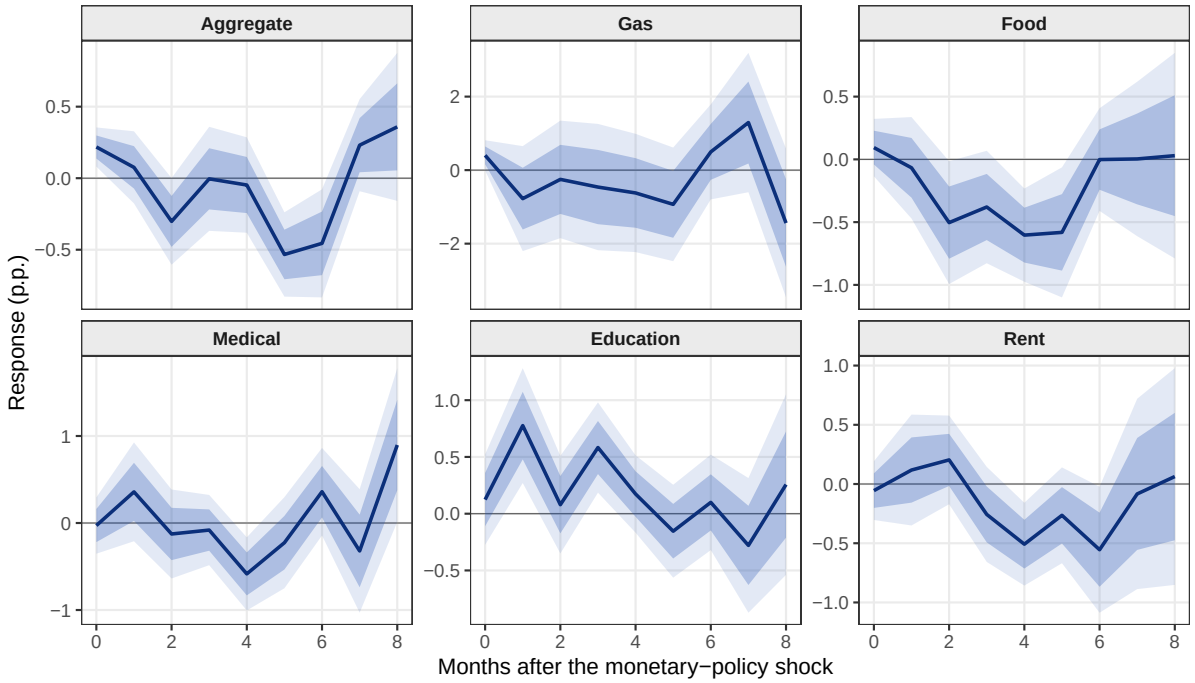


Figure 8: MP shock effect on sectoral price expectations

Note: This figure displays the estimated impulse response functions of household sectoral and aggregate price expectations to a 25 basis point monetary policy shock, identified using the Jarociński and Karadi (2020) approach. Impulse responses are estimated via local projections over an 8-month horizon. The outcome variables are aggregate and sector-specific one-year-ahead point forecasts from the Survey of Consumer Expectations (SCE). The five sectors are gasoline, food, rent, medical care, and education. Standard errors are clustered by survey month to account for the within-period correlation induced by the aggregate monetary policy shock, which is common to all respondents in a given month. Respondent fixed effects are included throughout. The sample runs from June 2013, with the shock series ending in January 2024; the acute COVID window (March–June 2020) is absorbed by outcome-date fixed effects rather than dropped (Appendix I). The projections also include outcome-date month-of-year fixed effects, which absorb category seasonality and remove the horizon-to-horizon wobble it induces. Shaded bands represent the 66.7% and 90% confidence intervals.

gories such as food and gasoline, where policy-induced price movements appear quickly in realized prices. In stickier categories, realized price dispersion should respond less directly to policy. We test this mechanism using BLS CPI micro-data. The exercise is closely related to evidence that monetary policy shocks generate heterogeneous inflation effects across households and consumption baskets (Lauper and Mangiante, 2025).

We construct price dispersion from detailed BLS CPI item indexes, using for each sector the most disaggregated items available under the relevant heading, mapped into six groups: food (91 items), gasoline (3 items), energy (7 items), rent (2 items), medical care (11 items), and education (5 items). For each item k in sector j and month t, we compute log price changes and define sectoral dispersion as the cross-sectional standard deviation of item-level inflation rates,

$$\text{Dispersion}_{j,t} = SD_{k \in j}(\pi_{k,t}) \,.$$

The baseline measure uses year-over-year item-level inflation, matching the annual horizon of the SCE inflation questions.

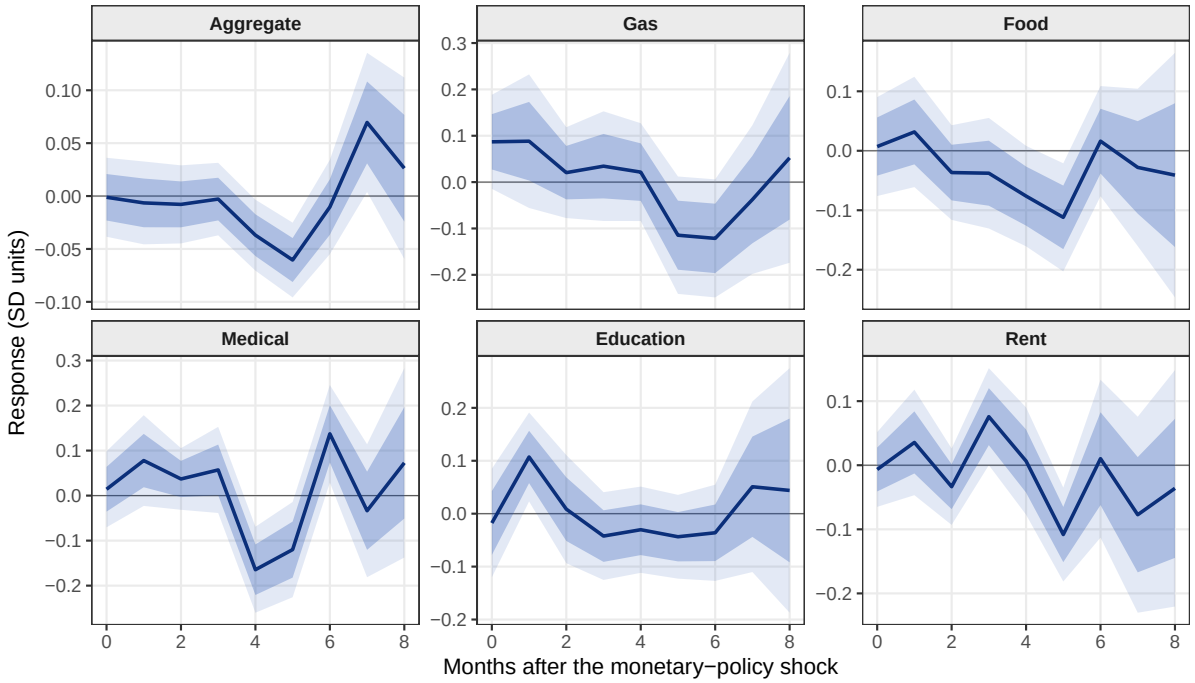


Figure 9: MP shock effect on sectoral price uncertainty

Note: Estimated impulse responses of household aggregate and sector-specific inflation uncertainty to a 25 basis point monetary policy shock. The outcome variables are the aggregate and sector-specific inflation-uncertainty indices from the SCE; specification, sample, COVID and seasonal fixed effects, clustering, and confidence bands are as in Figure 8.

Flexible versus sticky price sectors. Gasoline and food prices adjust frequently, while rent, medical care, and education adjust slowly because of long-term contracts, regulated prices, and annual tuition cycles. This difference in price flexibility provides the cross-sectional contrast we use to test the price-signal channel.

Figure 10 supports this mechanism. A contractionary shock compresses price dispersion significantly and persistently in the flexible food-and-energy sector, while the response of the sticky rent, medical, and education sector is muted.[14] The price-dispersion evidence links the survey responses in Figure 9 to a concrete price-based transmission mechanism: monetary policy reduces the dispersion of the item-level price signals households observe, especially in flexible categories, and this lowers the uncertainty attached to sectoral inflation. Appendix G.3 estimates this link directly in a sector-month panel with sector and month fixed effects; household uncertainty is significantly higher when the sector's realized price dispersion is wider.

Decomposing the response of non-core price dispersion clarifies the sources of the observed compression. Total cross-item variance in the non-core basket consists of two components: dispersion within the food category and a between-group component reflecting the squared difference between mean gasoline and food price changes (Appendix Figure J.13). These components respond differently to a contractionary monetary policy shock. The between-group component drives the significant impact response: the shock narrows the gas-food relative price gap during the first quarter, accounting for approx-

[14]The ECB reports similar results for the euro area: "Monetary policy pass-through to goods and services inflation: a granular perspective".

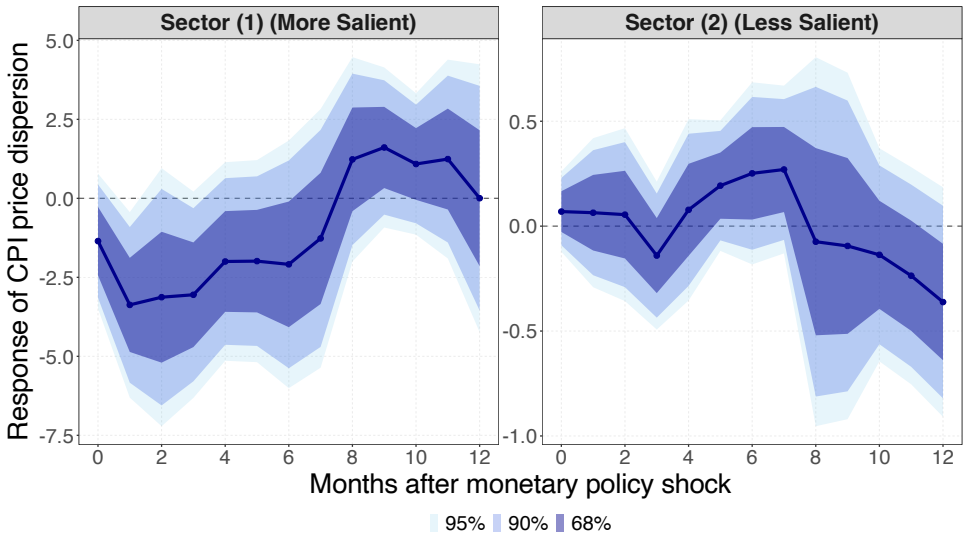


Figure 10: MP shock effect on sectoral price dispersion

Note: Impulse responses of sectoral CPI price dispersion (SD of YoY item-level inflation) to a 25bps Jarociński and Karadi (2020) monetary policy shock. Salient sector: food, gas, and energy (101 items). Less salient sector: rent, medical, and education (18 items). Controls: two lags of shock, outcome, and macro variables. Newey–West HAC standard errors with horizon-specific bandwidth. Shaded bands: 68%, 90%, 95% confidence intervals. Sample: June 2013–September 2024; the point estimates are essentially unchanged when March 2020 is excluded.

imately four-fifths of the total compression. By contrast, the within-food component changes little on impact and becomes positive in the third quarter, although the estimate is imprecise. This delayed increase is consistent with food retailers deferring price adjustments during the tightening and subsequently implementing them in a concentrated wave. The structural model incorporates both dynamics: movements in the gas–food relative price generate the immediate response of non-core dispersion, while deferred price adjustment produces its subsequent rebound (Section 6.2).

6 Should the Central Bank React to Goods-Specific Inflation Uncertainty?

The preceding empirical analysis established three key findings: (i) food, gas, and rent are the primary statistical contributors to household inflation uncertainty; (ii) uncertainty across salient items exhibits a strong covariance structure, clustering into non-durable necessities versus regulated services; and (iii) contractionary monetary policy reduces both the level and the uncertainty of inflation expectations, with the uncertainty decline building over four to six months, and compresses realized price-change dispersion almost exclusively in flexible food and gas prices. These findings raise the question of whether the central bank should respond to sectoral inflation uncertainty, and if so, to which sectors. Standard practice prescribes looking through volatile food and energy prices; we show that the opposite logic applies to the second moment. To make the argument we develop a structural noisy-information two-sector New Keynesian model in which the subjective posterior variance of sectoral inflation is endogenous to the policy stance. The mechanism is not assumed: because trend inflation is positive, monetary policy moves the cross-sectional dispersion of price changes at first order, and that dispersion

is the signal noise from which households infer inflation. Policy compresses dispersion far more in the flexible food-and-gasoline sector than in sticky services, so the central bank can materially move food-and-gas uncertainty but not services uncertainty. We solve the exact nonlinear Euler equation and the recursive Calvo reset-price block by pruned perturbation (Andreasen et al., 2018), an approach used to study the propagation of uncertainty in New Keynesian models (Oh, 2020); impulse responses are first-order and we evaluate welfare at the second-order pruned solution. Households and firms are Bayesian signal extractors throughout, so the belief block is micro-founded. At first order the perturbation solution coincides with the strict noisy-information equilibrium; at higher order it adds the precautionary and dispersion terms the welfare analysis uses. The welfare exercise then asks how a Taylor rule augmented with a response to sectoral posterior uncertainty should be set. Full derivations are in Appendix J.

6.1 A Structural Noisy-Information Two-Sector New Keynesian Model

The economy consists of two sectors, $j \in \{1, 2\}$, producing differentiated goods that are aggregated into a final consumption bundle with expenditure shares ω_1 and $\omega_2 = 1 - \omega_1$. Motivated by the within-individual correlation clusters of Section 4 (Figure 5b), we map Sector 1 to salient, non-durable necessities (food and gasoline) and Sector 2 to the rest of the economy (rent, medical care, and other services). Households have CRRA preferences and form posterior beliefs about sectoral inflation. The model counterpart of $\sigma^2_{\pi,j,t}$ is the subjective posterior variance of cumulative four-quarter sectoral inflation, matching the annual horizon of the SCE uncertainty measures; we suppress the horizon superscript below. Firms in each sector set prices subject to Calvo rigidities with non-adjustment probability α_j ($\alpha_1 = 0.70$ for food and gas, $\alpha_2 = 0.92$ for services, fixed to match the relative volatility of realized sectoral inflation; Table J.8), and the economy has positive trend inflation $\bar{\pi} = \log(1.02)/4$ per quarter (two percent annualized). The household and firm problems and their first-order conditions, the exact nonlinear consumption Euler equation (J.15) and the recursive Calvo reset-price block (J.21), are standard and are collected, with the complete equation system, in Appendix J.2; the welfare objects are derived in Appendix J.6. Five exogenous processes drive the economy: AR(1) sectoral markup-pressure states with weakly correlated innovations, an AR(1) common demand state, independent sectoral dispersion shocks, a services belief-noise shock, and the monetary policy shock; all are specified in Appendix J.2. The model deviates from this textbook economy in three places, each disciplined by the empirical results of Sections 4 and 5: positive trend inflation, which makes the cross-sectional dispersion of price changes respond to policy at first order; a noisy-information block in which that dispersion is the signal noise from which households infer category inflation, so posterior uncertainty inherits the dispersion's policy sensitivity; and belief loadings under which posterior means enter the allocations exactly where the latent states would under full information, so the posterior variance is a distinct object that policy can move without moving mean forecasts. Given this structure, we then test whether a Taylor rule that responds to non-core uncertainty improves welfare.

Each object of the Section 2 framework generalizes rather than disappears, and it is worth being explicit about the three substitutions the general-equilibrium version makes. The one-period signal about category inflation becomes a signal about the persistent latent

state driving it, so the static update of Equation (2) becomes a Kalman filter and the posterior variance inherits dynamics from the state; the exogenous signal variance $\sigma^2_{\eta,i,j,t}$ becomes the equilibrium price-change dispersion of the first deviation, which is what makes it policy-sensitive; and the household-specific objects, prior heterogeneity and cognitive noise, are not carried as structural states but live in the measurement layer, as the cross-sectional distribution around the common posterior and as the belief-noise residual whose scale matches the within-household correlation of the sectoral measures. The basket object of Equation (3) then reappears as U_t below, evaluated at the filter's posterior variances and correlation.

The first deviation is positive trend inflation, and its role is to turn price-change dispersion into a policy variable. The object that closes the model is the cross-sectional dispersion of price changes, which is what households sample when they form beliefs about category inflation. With Calvo pricing and positive trend inflation, the within-sector variance of reset-price gaps follows the recursion

$$\begin{aligned} \mu_{j,t} &= \alpha_j\Big(\mu_{j,t-1} - \pi_{j,t}\Big), \\ V_{j,t} &= \alpha_j V_{j,t-1} + \alpha_j(1-\alpha_j)\Big(\mu_{j,t-1} - \pi_{j,t}\Big)^2, \\ D^{\mathrm{YoY}}_{j,t} &= V_{j,t} + \Big(1 - 2\alpha_j^4\Big)V_{j,t-4}. \end{aligned} \tag{15}$$

where $\mu_{j,t}$ is the cross-sectional mean reset-price gap, carried as a state (not to be confused with the desired markup μ_j), and $D^{\mathrm{YoY}}_{j,t}$ is the variance of year-on-year sectoral price changes. In the linearization of (15) around the trend steady state, where $\bar{\mu}_j = -\alpha_j\bar{\pi}/(1-\alpha_j)$, sectoral inflation enters dispersion with a first-order impact loading $2\alpha_j\bar{\pi}$, which is nonzero precisely because $\bar{\pi} > 0$; at zero trend inflation dispersion responds to inflation only at second order (Ascari, 2004; Ascari and Sbordone, 2014). The mean-gap state $\mu_{j,t}$ then propagates the impulse forward, so the dispersion response builds over several quarters, with cumulative multiplier $2\alpha_j\bar{\pi}/(1-\alpha_j)$, rather than arriving at once. The impact loading is nearly common across sectors, so the sectoral asymmetry is carried by the size of the inflation response itself: the flexible food-and-gas sector ($\alpha_1 = 0.70$) has a much larger inflation response to policy than sticky services ($\alpha_2 = 0.92$), so its price-change dispersion, and hence the signal noise households face, is far more policy-sensitive, matching the cross-sector pattern of the estimated BLS dispersion responses (Figure 10).

The second deviation is the noisy-information block that reads this dispersion as household signal noise, the link documented directly in the sector-month regressions of Section 5.2, where household uncertainty is higher when a sector's realized price dispersion is wider. Writing $q_{j,t}$ for the log of the signal-noise variance,

$$q_{1,t} = \rho_q\, q_{1,t-1} + \kappa_{\mathrm{disp}}\,\hat{D}_{1,t}, \qquad q_{2,t} = \rho_q\, q_{2,t-1} + \kappa_{\mathrm{disp}}\,\hat{D}_{2,t} + \kappa_{\mathrm{serv}}\,\hat{D}_{1,t} + \sigma_{h,2}\,\varepsilon^h_{2,t}, \tag{16}$$

where $\hat{D}_{j,t} \equiv (D^{\mathrm{YoY}}_{j,t} - \bar{D}^{\mathrm{YoY}}_j)/\bar{D}^{\mathrm{YoY}}_j$ is the relative dispersion deviation, each sector reads its own dispersion with the common pass-through κ_{disp}, and services signal noise additionally reads the food-and-gas dispersion with the salience loading κ_{serv}: the most visible price dispersion households face contaminates their processing of all categories, the mechanism for which DAcunto et al. (2021) provide direct evidence in the level of beliefs, operating here on their precision. The memory ρ_q allows the perceived precision of beliefs to carry over across periods; we estimate it in Section 6.2 at 0.55, at which the model's quarterly

autocorrelation of food-and-gas uncertainty is 0.89 against 0.88 in the data. The residual $\sigma_{h,2}\varepsilon^h_{2,t}$ is idiosyncratic household uncertainty about services inflation, real belief noise rather than measurement error; we set its scale so that the within-household correlation of the two sectors' uncertainty measures matches the 0.34 measured in the SCE. Food-and-gas uncertainty carries no residual term, so it is entirely dispersion-driven, and the common-demand signal noise carries neither a dispersion loading nor, at baseline, a shock, so it is constant; the robustness analysis adds one and shuts the memory off ($\rho_q = 0$). A Bayesian updating block maps the signal-noise states $q_{j,t}$ and the latent sectoral markup and common-demand states into the posterior means $m_{j,t}, m_{d,t}$ and the posterior variances $\sigma^2_{\pi,j,t}$ and correlation ρ_t (Appendix J.5); we implement the filter in square-root (Cholesky) form so the posterior covariance is positive-definite at every perturbation order. There is no free loading of policy into signal noise: both dispersion inputs are equilibrium objects, κ_{disp} and κ_{serv} are estimated in Section 6.2 jointly with the other transmission parameters, and the cross-sector ordering of uncertainty responses is anchored by the primitive stickiness $\alpha_1 < \alpha_2$, which makes the food-and-gas dispersion the dominant, policy-sensitive input on both lines of (16).

This dispersion channel is the general-equilibrium counterpart of the signal-extraction framework of Section 2. A household's signal about category-j inflation is a noisily sampled item-level price change, so the cross-sectional dispersion of item price changes within the category is the variance of the signal noise (the $\sigma^2_{\eta,i,j,t}$ of Equation (1)). The unconditional relation between inflation and price dispersion is contested: Nakamura et al. (2018) find no rise in price dispersion during the Great Inflation, and Sheremirov (2020) documents a weak unconditional link. The mechanism here is disciplined instead by the conditional response of dispersion to an identified monetary shock, estimated directly in Figure 10. The signal-extraction block is therefore structural: policy moves dispersion, dispersion is the noise in households' price signals, and Bayesian updating turns that noise into posterior uncertainty. The one behavioural object this block adds is κ_{disp}, the intensity with which dispersion passes into signal noise. It is an estimated pass-through rather than a structural elasticity: a number that stands in for households resolving only a limited number of the item prices in a category, so the dispersion they cannot resolve becomes belief uncertainty (Sims, 2003; Maćkowiak et al., 2023). The salience loading κ_{serv} formalizes the attention spillover introduced in Section 2: attention is a household-level resource, and the dominant dispersion households face, food and gasoline, occupies part of the capacity they would otherwise use to resolve services prices. The sign and gradual build of the services uncertainty response then follow from the food-and-gas dispersion dynamics, with no timing parameter of their own.

The BLS dispersion of Section 5.2 is an item-level statistic with two moving parts, dispersion within the food group and the gap between the gas and food group means (Appendix Figure J.13), and the model carries both: the between component is driven by a gas–food relative price whose energy-demand state adjusts slowly through the vehicle stock (Hughes et al., 2008), and measured dispersion additionally reflects price adjustments deferred during a tightening and released in a bunched wave afterwards, the selection dynamics of state-dependent pricing (Caballero and Engel, 2007; Golosov and Lucas, 2007). These additions live entirely in the measurement layer: they change no allocation, pricing decision, or welfare object, and the uncertainty calibration is numerically identical with and without them (Appendix J.4).

The third deviation determines where beliefs enter, and it is what gives the survey measures a model counterpart: the SCE objects of Section 3 are subjective posterior variances, and $\sigma^2_{\pi,j,t}$ is their model analogue, produced by a square-root Kalman block that is the dynamic, multi-signal counterpart of the static update in Equation (2). Households and firms are Bayesian signal extractors over the latent sectoral markup and common demand states, and their beliefs reach the allocations in exactly two places, both standard first-order conditions shifted by a posterior mean. The household's Euler equation carries an effective preference shifter $\xi_t^{\text{eff}} \equiv \exp(\log \xi_t + \lambda_d \, m_{d,t})$, so beliefs about latent demand move consumption (Appendix Equation (J.15)); the firm's reset-price problem carries an effective marginal cost shifted by the posterior mean $m_{j,t}$ of sectoral markup pressure (Appendix Equation (J.21)). We set both loadings to one: the posterior mean enters exactly where the latent state would enter under full information, so imperfect information is the only new primitive. When the rule moves the policy rate, the movement propagates to consumption through the Euler equation, with elasticity $1/\sigma_c$, and on to sectoral inflation through the reset-price block, with the standard Calvo slope $\kappa_j = (1-\alpha_j)(1-\alpha_j\beta)/\alpha_j$ (distinct from the estimated dispersion pass-throughs $\kappa_{\text{disp}}, \kappa_{\text{serv}}$); the induced movement in sectoral inflation then feeds the dispersion recursion of Departure 1, closing the loop on next-period uncertainty. The posterior *variances*, by contrast, enter no allocation; they enter only the policy rule. The aggregate posterior basket-uncertainty object is

$$U_t \equiv \omega_1^2 \sigma^2_{\pi,1,t} + (1-\omega_1)^2 \sigma^2_{\pi,2,t} + 2\omega_1(1-\omega_1)\rho_t \sigma_{\pi,1,t}\sigma_{\pi,2,t} - \bar{U}, \tag{17}$$

measured as a deviation from its steady-state value $\bar{U}$ and mirroring the variance decomposition in Equation (3); it depends on the posterior conditional correlation ρ_t between sectoral inflation forecasts. Policy moves the signal-noise variance, a second moment: at first order the Kalman gains sit at their steady-state values and the signal realizations are unaffected, so the subjective mean forecast of sectoral inflation does not respond to a policy shock at first order; the stance is, to first order, a pure inflation-uncertainty instrument, distinct from the standard level channel.[15]

These three deviations complete the structure. The policy experiment then asks whether, in an economy with this empirically disciplined belief block, a rule that responds to the resulting uncertainty objects improves welfare. Define the normalized sector-j uncertainty gap $h_{j,t} \equiv \log \sigma^2_{\pi,j,t} - \log \bar{\sigma}^2_{\pi,j}$, the log deviation of posterior sectoral uncertainty from its steady state. The central bank follows a Taylor-type rule with two sectoral margins, one on the price *level* and one on *uncertainty*:

$$i_t = \rho_i \, i_{t-1} + (1-\rho_i)\Big(\bar{\imath} + \phi_\pi\Big(\theta_1(\pi_{1,t} - \bar{\pi}) + (1-\theta_1)(\pi_{2,t} - \bar{\pi})\Big) + \phi_x(x_t - \bar{x}) + \phi_{U,1} \, h_{1,t} + \phi_{U,2} \, h_{2,t}\Big) + \sigma_i \, \varepsilon_t^i. \tag{18}$$

Here $\bar{\imath}$ and $\bar{x}$ are the trend steady-state nominal rate and output, ε_t^i is the monetary policy shock, and $\sigma_i = 0.0025$ (25 basis points per quarter); since the estimation matches proportional responses, this normalization is innocuous, and we report impulse responses per 25 basis points to match the empirical local projections. The conventional coefficients

[15] At higher order the policy-driven signal noise interacts with the signal innovations, so the mean-forecast response is not exactly zero; numerically it is negligible relative to the uncertainty response, and we report the zero response as a first-order property.

are held fixed at $\rho_i = 0.3$, $\phi_\pi = 2$, and $\phi_x = 0$ ($\phi_x > 0$ is destabilizing under trend inflation, Appendix J.15; Table J.10 shows the conclusions are unchanged at $\rho_i = 0$ and $\rho_i = 0.7$), while the two sectoral margins $(\theta_1, \phi_{U,1}, \phi_{U,2})$ are the choice variables we optimize. The inflation term is a sectoral-weighted gap, with weight θ_1 on food-and-gas inflation and $1 - \theta_1$ on services inflation (the rule weight θ_1 is unrelated to the within-sector demand elasticity θ_j of Appendix J.2): $\theta_1 = \omega$ is the expenditure-weighted (headline) rule and $\theta_1 = 0$ is core (services) targeting, which looks through the food-and-gas level. The coefficients $\phi_{U,j}$ measure the response to a one-log-point deviation of sector-j posterior uncertainty. The policy question has two parts that we optimize jointly: how much weight the level margin should place on food-and-gas inflation (θ_1), and whether the uncertainty margin should lean against food-and-gas uncertainty ($\phi_{U,1} > 0$) and services uncertainty ($\phi_{U,2} > 0$) or look through them. Section 6.3 shows that the optimum sets $\theta_1^* = 0$ and $\phi_{U,1}^* > 0$: the rule looks through the food-and-gas level but leans against its uncertainty, and does the reverse for services. We rank rules by the representative household's exact expected lifetime utility in consumption and hours (Appendix J.6) rather than by a quadratic proxy. Because the model carries sectoral price-dispersion states, welfare internalizes misallocation through hours worked directly. We evaluate welfare at the true state of the economy: the felicity weight is priced at the true common-demand state $z_{d,t}$ while households continue to choose consumption and labour under their noisy posterior beliefs $m_{d,t}$, so the criterion is the true lifetime utility of an allocation formed under imperfect information rather than the utility agents mistakenly believe they attain.[16]

The model connects four literatures. It shares the information-friction mechanism of the rational-inattention and incomplete-information New Keynesian models (Sims, 2003; Mankiw and Reis, 2002; Lorenzoni, 2009; Nimark, 2008; Maćkowiak et al., 2023; Coibion and Gorodnichenko, 2015a), but makes signal noise itself respond to the policy stance through the equilibrium dispersion of price changes rather than through an assumed loading, so policy moves the subjective variance of beliefs and not only their level. The closest normative antecedent is Paciello and Wiederholt (2014), who characterize optimal policy when firms' information is endogenous to policy; here the policy-sensitive object is instead the posterior variance of sectoral inflation, disciplined by a measured survey counterpart, and the endogeneity runs through the price-change dispersion that Calvo staggering generates under positive trend inflation. Melosi (2017) offers a complementary signalling channel in which policy actions themselves reveal information; the transmission documented here runs through realized prices rather than announcements. Ascari et al. (2023) study the amplification from endogenous uncertainty after shocks to inflation expectations; we make the uncertainty policy-sensitive and ask how the rule should respond to it. It adopts the two-sector optimal-policy structure of the sectoral-rigidity literature (Aoki, 2001; Benigno, 2004; Erceg and Levin, 2006; Mankiw and Reis, 2003), but the sectors differ only in their frequency of price adjustment ($\alpha_1 < \alpha_2$), and it is this primitive that governs how strongly each sector's uncertainty responds to policy, reaching a prescription that leans on the flexible food-and-gasoline sector rather than on the stickiest one. And it operationalizes the welfare-cost-of-uncertainty tradition (Friedman, 1977; Cukierman and Meltzer, 1986; Ball, 1992) inside an estimable environment

[16] Since $\mathrm{Var}(z_{d,t}) > \mathrm{Var}(m_{d,t})$, evaluating at the true state rather than the perceived one could shift the measured cost of uncertainty; in practice the two objects agree to four significant figures, because the common-demand belief contributes negligibly relative to the sectoral dispersion channel. We report the true-state object throughout and the perceived-state object in Appendix J.6.

with a policy-responsive posterior uncertainty object. Three features are new relative to existing noisy-information models: signal noise is tied to the model's own equilibrium price-change dispersion; a cross-category salience loading lets the most visible food-and-gas dispersion enter services signal noise, so the muted services response is reproduced with the right timing, which no parameter fits, as an over-identifying restriction; and the Taylor rule responds to sector-specific posterior uncertainty gaps rather than to a single aggregate index. Appendix J.1 develops these connections, along with the firm-side endogenous-slope complement of Flynn et al. (2026).

6.2 Estimating the Transmission Parameters

The central transmission parameters are the pass-through κ_{disp} from equilibrium price-change dispersion to signal noise and the salience loading κ_{serv} in Equation (16). The pass-through is the intensity of the model's one information friction: because households process only limited information about the many item prices in a category (Sims, 2003; Maćkowiak et al., 2023), the cross-sectional dispersion of price changes they cannot fully resolve maps into the noise in their inflation signal. It also doubles as the mapping between the model's posterior inflation variance and the survey-based uncertainty measure it is matched to. The policy sensitivity of the driving input is pinned by primitives: the trend-inflation loading in Equation (15) and the sectoral stickiness $\alpha_1 < \alpha_2$. We estimate κ_{disp}, the salience loading κ_{serv}, the belief memory ρ_q, and the parameters of the measured-dispersion block jointly by minimum distance, in the impulse-response-matching tradition, holding the externally identified parameters fixed: the sectoral rigidities, the policy-rule coefficients, trend inflation, preferences, the survey underreaction rigidity, and the measurement constants that are direct data ratios. Two further constants are measured rather than estimated, each against one moment: the cross-sector coupling of the markup innovations, $\text{corr}_z = 0.055$, reproduces the near-zero correlation of realized sectoral inflation (0.04 in the data), and the idiosyncratic services-uncertainty scale $\sigma_{h,2}$ reproduces the within-household correlation of the two sectors' uncertainty measures (0.34). The food-and-gas moments identify κ_{disp} and the services moments identify κ_{serv}, and the two estimates are essentially uncorrelated; what remains over-identifying about the services response is its timing and shape, which have no parameter of their own.

The estimator minimizes the standard-error-weighted distance between model and data over nineteen moments: the sectoral uncertainty impulse responses (the covariance-aware respondent-level composites of Section 5, averaged to the model frequency), the between and within components of the non-core dispersion response, the gas–food relative-price response scored on shape, and the persistence of the measured uncertainty index; Appendix J.13 states the full moment set and weights. The model's deterministic gas–food channel cannot, however, reproduce the sign pattern of the empirical relative-price response at any positive scale, so its shape scale floors at zero and the four gap moments act as a constant over-identifying penalty that no parameter can reduce. Scales are concentrated out by weighted least squares, so only proportional responses are matched, consistent with the ordinal use of the survey measure in Section 3. We run the estimator from dispersed starting values; the reported point attains the best fit and is reproduced by multiple starts. At the estimate, $\kappa_{\text{disp}} = 2.08$, $\kappa_{\text{serv}} = 0.83$, and $\rho_q = 0.55$, the model's cumulative food-and-gas uncertainty response is -9.9 against -9.2 in the data, within the empirical 90% band at every quarter, building gradually to its trough as the empirical response does; the persistence of measured uncertainty is 0.89 against 0.88; and the be-

tween component carries the impact compression of dispersion, as in the decomposition of Appendix Figure J.13. The estimates carry wide standard errors (2.0 for κ_{disp}, 1.7 for κ_{serv}, and 0.22 for ρ_q, from the bootstrap moment covariance by the delta method), as the local projections they fit are themselves borderline significant; the policy conclusions do not hinge on the point, since the robustness battery of Appendix J.14 re-optimizes the rule at pass-through values from 1.0 to above the estimate and the leaning survives throughout, with larger gains at smaller pass-through. Appendix J.13 states the estimator, the moment set, the constrained parameters, and the inference.

The size of the services (Sector 2) response is fitted, through κ_{serv}, but its timing and shape are not: the salience channel imports the food-and-gas dispersion dynamics with no timing parameter, and the model's services trough arrives in the third quarter, as in the data, at -1.9 against an empirical trough-quarter response of -2.3 whose 90% interval $[-10.0, 5.4]$ includes zero, with a services-to-food-and-gas trough ratio of 0.42 against 0.40. Both local projections lose significance beyond the trough; where the empirical response is insignificant we treat the effect as having decayed to zero, and the model reproduces the gradual decline in both sectors, smaller in services than in food and gas, as in the data.

We estimate the measured-dispersion block within the same system, subject to the measured constants: the unconditional gap volatility (18pp) is imposed as an identity, the between weight is the measured between-to-within variance ratio (0.757), and the remaining parameters are estimated, with three at economically meaningful bounds reported as constrained: the activity share of the gap at its floor of zero, the catch-up share at its intensity cap, and the gap-state persistence at its floor; the policy sensitivity of oil-innovation volatility is an interior estimate. Substantively, the data attribute the policy-conditional movement of the gas–food gap to the volatility margin rather than the activity channel, and profiling the activity share confirms zero is preferred. With the estimates in place the model reproduces the shape of the measured non-core dispersion response in Figure 11: compression on impact, matched almost exactly, and the crossing into positive territory as deferred adjustments release; the core dispersion response remains a near-zero prediction. Because the between and deferral components are quadratic and anticipated objects, the plotted dispersion response is the model's mean response at second order; Appendix J.4 details the block and the verification that the uncertainty moments and the welfare results are numerically identical with the block removed.

Table 6 summarizes the estimated moments, the over-identifying restrictions, and the untargeted diagnostics. The size of the services response is fitted; the over-identifying content is its unparameterized timing, the near-orthogonality of uncertainty and inflation, and the untargeted realized-inflation moments. Appendix J.10 discusses the fit and plots the model impulse responses against the empirical local projections (Figure J.14). The remaining parameters follow standard quarterly New Keynesian calibration and are reported in Table J.8.

The first moment of beliefs reacts through two channels. The first is standard: the policy shock moves the model's sectoral inflation and hence households' forward-looking expectations, and reported forecasts underreact with well-documented information rigidity, so the expectations rows of Figure 11 map the model's inflation path through the measurement equation of Coibion and Gorodnichenko (2015a) at their estimated rigidity ($\lambda = 0.55$, not fitted). The second channel runs through uncertainty itself: households

Table 6: Model Fit: Estimated Moments, Over-Identifying Restrictions, and Untargeted Moments

Moment	Data	Model	Parameter
A. Estimated moments *(minimum distance; parameter point estimates)*			
Food/gas uncertainty IRF, cumulative 3Q (%)	−9.2	−9.9	$\kappa_{\text{disp}} = 2.08$
Food/gas uncertainty, quarterly AR(1)	0.88	0.89	$\rho_q = 0.55$
Services uncertainty IRF, trough (%)	−2.3	−1.9	$\kappa_{\text{serv}} = 0.83$ (salience loading)
Dispersion between-component, impact (var. units)	−25.4	−23.9	ψ_v (volatility margin)
Dispersion within-component, Q3 rebound	+9.3	+3.5	ρ_b, ν (deferral–catch-up)
Cross-sector corr. of realized inflation	0.04	0.04	$\text{corr}_z = 0.055$ (measured)
Within-household corr. of sectoral uncertainty	0.34	0.34	$\sigma_{h,2} = 5.25$ (measured)
Realized inflation-volatility ratio, S1 / S2	6.1	10.8	$\alpha_1, \alpha_2 = 0.70, 0.92$
B. Over-identifying restrictions *(no dedicated parameter; tested against data)*			
Services uncertainty trough timing (quarter)	Q2–3	Q3	— (no timing param.)
Services / food-gas trough ratio	0.40	0.42	— (follows from fit)
C. Untargeted moments *(reported for transparency, not fit)*			
Realized 4Q inflation SD, food/gas (pp)	4.8	3.0	—
Realized 4Q inflation SD, services (pp)	0.8	0.28	—
Correlation of food/gas uncertainty and inflation	—	−0.11	—

Notes: **Panel A** lists the moments entering the minimum-distance estimation (Appendix J.13); the empirical uncertainty targets are local projections on respondent-level expenditure-weighted sector composites, so cross-category covariances are embedded, averaged to quarters. The persistence moment is matched at the estimate; the between and within dispersion components are fitted under one common scale, and the listed parameters are those most strongly identified by each moment. (α_1, α_2) are held fixed, chosen to match the relative volatility of realized sectoral inflation, consistent with regular-price frequency evidence in Nakamura and Steinsson (2008).

Panel B reports over-identifying restrictions. The size of the services response is fitted (Panel A), but its timing and shape carry no parameter: the salience channel imports the food-and-gas dispersion dynamics, so the third-quarter services trough and the gradual build are predictions. The empirical services confidence band includes zero, so the size row in Panel A is disciplined loosely and the timing is the sharper test.

Panel C reports untargeted moments. The model understates the levels of realized four-quarter inflation volatility but matches the near-zero cross-sector correlation of realized inflation and produces near-zero comovement between food/gas uncertainty and inflation. Model values are pruned ergodic moments and impulse responses of the nonlinear two-sector model at the baseline calibration $\rho_i = 0.3$, $\phi_\pi = 2$, $\phi_x = 0$, and $\sigma_i = 0.0025$, under a rule that does not respond to uncertainty ($\phi_{U,1} = \phi_{U,2} = 0$). A homogeneous rescaling of the fundamental and signal-noise volatilities that matches the food-and-gas level leaves the estimation exactly invariant and raises the welfare magnitudes about threefold (Appendix J.14). Empirical moments are from the FRBNY SCE micro panel and BLS CPI item indices.

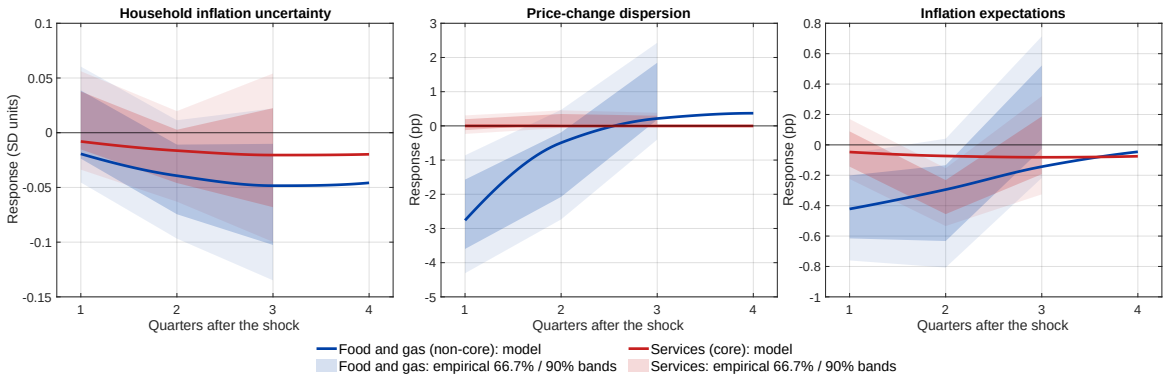


Figure 11: Model Fit: Targeted and Untargeted Responses

Note: Each panel overlays both sectors on a shared vertical scale: food and gas (non-core) in blue and services (core) in red. Shaded regions are the empirical 66.7% (dark) and 90% (light) local-projection bands, aggregated to quarters one through three (quarterly standard errors aggregate the monthly ones assuming independence across months); the central empirical estimate is omitted so that the model response can be read against the bands directly. Solid coloured lines are the model responses, drawn one quarter beyond the horizon the projections support. The three panels are the observables of the mechanism: household inflation uncertainty (left) and inflation expectations (right) are the covariance-aware sector composites from the SCE, and price-change dispersion (middle) is the BLS item-level dispersion projection of Section 5.2, whose non-core model line is the measured-dispersion object of Appendix J.4 and crosses zero in the third quarter as in the data; the expectations panel passes the model's inflation path through the underreaction measurement equation of Coibion and Gorodnichenko (2015a) at $\lambda = 0.55$ and adds the measured uncertainty premium of reported forecasts (5.6 and 6.7 percentage points per unit of the respective χ-based uncertainty measure, times the model uncertainty response; the premium is in data units and is not rescaled). Within each panel a single scale is fitted to the non-core sector by least squares, so the core response is a prediction with no free normalization. On the shared scale the muted core response is directly visible against the larger non-core one. Model lines are shape-preserving interpolations of the quarterly responses. The acute COVID window (March–June 2020) is absorbed by outcome-date fixed effects, dropping no observation (Appendix I).

reporting more uncertainty report systematically higher point forecasts, the mean of a right-skewed subjective distribution. Within respondent and month, a unit of the zeta uncertainty measure raises the reported category forecast by 5.6 percentage points for the food-and-gas composite and 6.7 for services (cluster t-statistics near 100; measured, not fitted to any impulse response), and the causal counterpart is the randomized finding that resolving households' uncertainty lowers their reported expectations (Fischer et al., 2025). The expectations rows therefore add this uncertainty premium, the measured coefficient times the model's uncertainty response, to the CG-filtered path: when policy compresses uncertainty, reported expectations fall by more than the inflation path alone implies, at exactly the quarters where the uncertainty troughs sit. The premium roughly doubles the model's core expectations response, to -0.06 in the second quarter against -0.33 in the data, covering about a fifth of the dip; the composite's own components show what carries the rest: rent has 72 percent of the core composite weight and its expectations response is the only individually significant one (-0.34, $t = -2.2$), the housing channel of monetary transmission, which the model's services sector, built to carry the uncertainty margin rather than housing, does not contain. The inflation-uncertainty channel operates alongside the conventional expectations response of monetary policy, not in place of it.

6.3 Optimal Policy: Look Through the Food-and-Gas Level, Lean on Its Uncertainty

We now ask how the central bank should set its response to sectoral inflation and to sectoral uncertainty, ranking rules by the household's exact expected lifetime utility evaluated at the true state of the economy. The optimal rule inverts the textbook prescription across moments. For the price level, the model reproduces the standard advice: look through volatile food and energy prices and target core inflation. For the uncertainty about those same prices, the optimal response is the opposite: lean against food-and-gas uncertainty and look through services uncertainty. Both parts follow from price stickiness. Misallocation from the price level originates in the sticky sector, so that is the level the bank stabilizes; the dispersion that policy can compress originates in the flexible sector, so that is the uncertainty the bank leans against.

For this contrast to be more than a relabeling, sectoral uncertainty must be able to move without the sectoral price level moving with it; otherwise responding to uncertainty and responding to inflation would be the same instrument. In the model they are not, and this is a structural consequence of imperfect information rather than an added assumption. Sector-j uncertainty $h_{j,t} = \log \sigma^2_{\pi,j,t}$ is the posterior variance of a signal-extraction problem: households observe price changes with imperfect precision, so their belief uncertainty responds to the noise in their signals, which is only partly driven by realized sectoral inflation. At the estimated calibration, under a rule that does not respond to uncertainty, the correlation between food-and-gas uncertainty and food-and-gas inflation is about -0.11: because the dispersion that drives uncertainty accumulates through the mean-gap state rather than tracking current inflation, the two are essentially orthogonal objects, and a rule can respond to one without mechanically responding to the other. Appendix J.14 confirms the results are unchanged when we add an explicit dispersion shock: idiosyncratic shifts in the cross-sectional spread of price changes that leave the sectoral mean unchanged and, at first order, move the posterior variance without moving inflation.

The rule has two margins: the level weight θ_1 on food-and-gas inflation, where $\theta_1 = \omega = 0.27$ is the headline rule and $\theta_1 = 0$ is core targeting, and the food-and-gas uncertainty response $\phi_{U,1}$. Figure 12 plots the welfare loss over these two margins. The joint welfare-maximizing rule puts zero weight on food-and-gas inflation, $\theta_1^* = 0$, and a positive weight on food-and-gas uncertainty, $\phi_{U,1}^* = 0.003$, a sharp interior optimum. Reading the figure down any column, welfare falls as the bank puts more weight on the food-and-gas price level; even the expenditure share is too much weight. Reading across the optimal row, welfare falls steeply on both sides of the optimum, more steeply below it: full look-through of uncertainty ($\phi_{U,1} = 0$) is welfare-dominated by about 25 basis points of consumption, and the conventional headline rule ($\theta_1 = \omega$, no uncertainty response) by about 35 basis points of consumption every period; both magnitudes are at the conservative baseline volatility scale, and the level-matched rescaling of Appendix J.14 roughly triples them. These are large stakes: classic estimates put the welfare cost of all business-cycle consumption volatility at about 5 basis points (Lucas, 2003). The optimal rule looks through the food-and-gas level and leans against food-and-gas uncertainty at the same time. These magnitudes are conservative with respect to the model's volatility scale: realized food-and-gas inflation volatility sits below the data (3.0 against 4.8 percentage points), and the homogeneous volatility rescaling that closes the gap leaves

the estimation exactly invariant while raising the consumption-equivalent stakes about threefold, with the optimum unchanged (Appendix J.14).

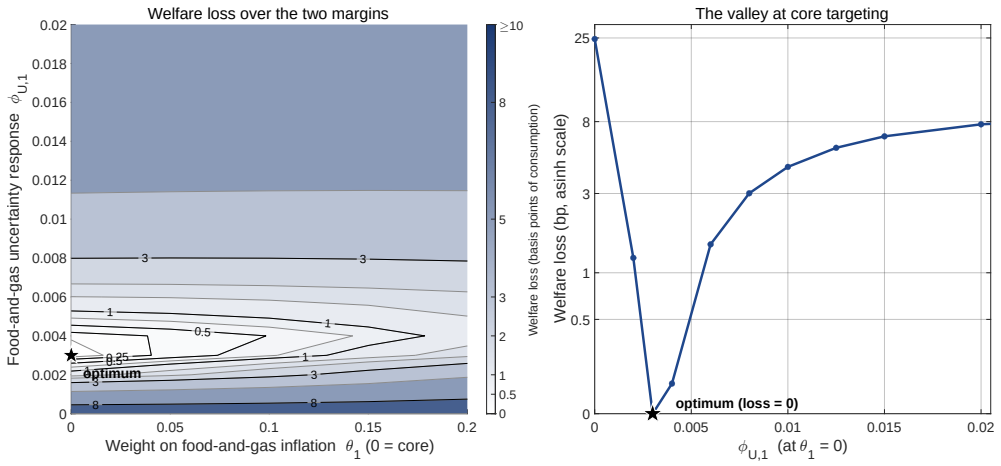


Figure 12: The Inversion: Look Through the Food-and-Gas Level, Lean on Its Uncertainty

Note: Exact-utility welfare loss in basis points of consumption, relative to the welfare-maximizing rule, evaluated at the second-order pruned solution, over the weight on food-and-gas inflation in the Taylor rule (horizontal; 0 = core targeting; the display is truncated where the loss saturates, and the headline weight $\omega = 0.27$ lies beyond the right edge, deeper in the saturated region) and the food-and-gas uncertainty response $\phi_{U,1}$ (vertical), at the estimated calibration ($\phi_\pi = 2$, $\phi_x = 0$, $\rho_i = 0.3$, $\kappa_{\text{disp}} = 2.08$, $\bar{\pi} = \log(1.02)/4$). Welfare is evaluated at the true state while agents choose under noisy beliefs; posterior uncertainty varies essentially orthogonally to the price level (correlation -0.11) because it is the variance of a signal-extraction problem driven by accumulated dispersion. The left panel plots the loss over the two margins; labeled contours mark losses of 0.25, 0.5, 1, 3, and 8 basis points, and losses above 10 basis points are shown in the darkest shade so that the colour scale resolves the valley around the optimum. The right panel plots the same loss along the core-targeting column ($\theta_1 = 0$) on an inverse-hyperbolic-sine scale, which compresses large losses like a log axis while displaying the exact zero at the optimum: the loss sits at about 25 basis points when the rule ignores uncertainty and falls steeply to the interior optimum at $\phi^*_{U,1} = 0.003$. The valley is asymmetric: under-responding to uncertainty is substantially costlier than over-responding, and losses mount toward the headline rule in the bottom-right corner. The star marks the joint optimum.

The two margins are governed by the same sectoral stickiness in opposite ways. For the level, the logic is Aoki (2001): flexible food-and-gas prices ($\alpha_1 = 0.70$) adjust efficiently and generate little relative-price misallocation, so stabilizing their level buys almost nothing, while sticky services prices ($\alpha_2 = 0.92$) do generate misallocation and are worth stabilizing. For the uncertainty, the logic runs through the dispersion channel of Section 6.1: policy compresses price-change dispersion, and hence signal noise, only where prices actually move, which is the flexible sector. The bank can materially reduce food-and-gas uncertainty but not services uncertainty, so the uncertainty response concentrates where policy has traction. The same $\alpha_1 < \alpha_2$ that makes food and gas the wrong target for the level makes it the right target for the uncertainty. Figure 13 makes the contrast a fair one by writing each margin as a food-and-gas share of a fixed reaction budget, reallocated against services, so looking through one sector always means leaning on the other. The optimal food-and-gas share is zero for the inflation level and one for the uncertainty. The mirror image holds on the services side: optimizing the two sectoral uncertainty responses jointly sets the services response to zero, $\phi^*_{U,2} = 0$, because services prices adjust infrequently, so policy moves services price-change dispersion very little and a response to it buys negligible stabilization while adding interest-rate volatility.

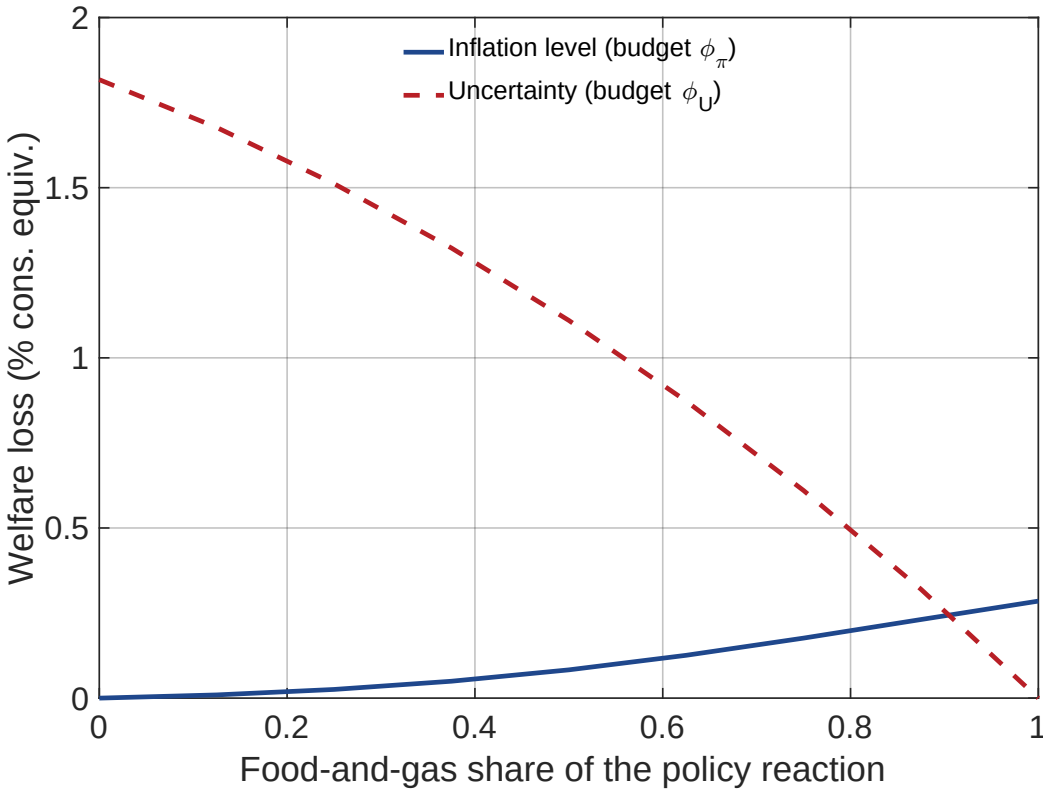


Figure 13: The Optimal Food-and-Gas Share of the Policy Reaction: Inflation Level versus Uncertainty

Note: Each line asks the same question for one margin of the rule: holding that margin's total reaction fixed, what fraction should fall on food and gas rather than on services? The horizontal axis is that fraction; the vertical axis is the welfare loss (consumption-equivalent, relative to the joint optimum). The solid line reallocates the inflation response ϕ_π: the loss is smallest at share zero, so the level response belongs entirely on services (core targeting). The dashed line reallocates the uncertainty response ϕ_U: the loss is smallest at share one, so the uncertainty response belongs entirely on food and gas. Because each budget is fixed, looking through one sector on a margin necessarily means leaning on the other.

The optimal food-and-gas uncertainty coefficient, $\phi^*_{U,1} = 0.003$, is small next to the inflation coefficient $\phi_\pi = 2$, but the two are not on the same scale: the uncertainty term multiplies the log posterior variance h_1, whose standard deviation (about 1.2 in log-variance units under the optimal rule) is far larger than that of the core inflation gap. In basis points per standard deviation the response is substantial: a one-standard-deviation rise in food-and-gas posterior uncertainty raises the annualized policy rate by about 100 basis points through the rule's uncertainty term, whereas a one-standard-deviation rise in the targeted core inflation gap raises it by about 50 basis points through ϕ_π. In rate terms the bank therefore leans about twice as hard on food-and-gas uncertainty as on core inflation, and the uncertainty term accounts for most of the variance of the systematic component of the policy rate (86 percent). The uncertainty in the rule is not manufactured by policy: households' imperfect information and exogenous food-and-gas disturbances account for almost all of the variance of aggregate posterior uncertainty, while the monetary policy shock itself accounts for a negligible share. The rule asks the bank to lean against uncertainty driven by non-core supply and by the noise in households' information, the same sources at issue in the look-through debate.

The rule is determinate over the empirically relevant range: the sectoral rule satisfies

the Blanchard–Kahn conditions for $\phi_{U,1}$ up to at least 0.1, more than thirty times the optimum, and for $\phi_{U,2}$ up to at least 0.02. Under positive trend inflation the determinate region is smaller than in the zero-inflation benchmark and is sensitive to the response to the output gap; the baseline sets no separate response to the output gap ($\phi_x = 0$), which keeps the economy determinate at a conventional inflation coefficient $\phi_\pi = 2$. Appendix J.15 reports the determinacy diagnostics for the sectoral rule.

6.4 Discussion

Food and gas prices are important contributors to household inflation uncertainty (Section 4) and the categories in which contractionary monetary policy most effectively compresses price-change dispersion and uncertainty (Section 5). The model therefore prescribes a "risk-management" approach: the central bank should lean against food-and-gas uncertainty while looking through the corresponding price-level movements and targeting core inflation.

This prescription is distinct from a conventional response to an inflationary shock. In the baseline simulations (Appendix J.11), an uncertainty shock can lower aggregate inflation on impact because the contraction in real activity dominates the price response. The optimal tightening is instead a response to the uncertainty state itself: it reduces signal noise and future inflation variance, thereby limiting the welfare costs associated with output-gap and price-dispersion dynamics. The channel decomposition in (Appendix J.7) attributes approximately 97 percent of these costs to firms' mispriced Calvo resets and the resulting relative-price misallocation. Policy leans against food-and-gas uncertainty because its dispersion-driven component is compressible; services uncertainty is largely driven by components over which policy has little traction.

The result is robust to whether firms share households' uncertainty, to a salience-weighted perceived consumption basket, and to dispersion shocks that move uncertainty independently of the price level (Appendix J.14).

The rule is implementable using the monthly survey-based sectoral uncertainty indices constructed in Section 3, rather than an unobservable model state. At the optimum, the annualized standard deviation of the policy rate rises from approximately 120 to 175 basis points, while aggregate inflation volatility falls from approximately 135 to 90 basis points. Goods-specific inflation uncertainty is therefore worth monitoring alongside aggregate measures, but the optimal response remains moderate and state-contingent, balancing lower posterior uncertainty and price-dispersion losses against greater interest-rate volatility.

7 Conclusion

We develop sector-specific measures of subjective inflation uncertainty, identify which sectors drive it, and draw out what it implies for monetary policy. Extending the round-numbers approach of Binder (2017b) to category-level price forecasts in the SCE, we construct and validate monthly uncertainty indices for gasoline, food, medical care, education, and rent from 2013 to 2025. Food-price uncertainty accounts for the largest share of the explained variation in aggregate inflation uncertainty, followed by rent, education, and gasoline, whose contribution increases after 2020.

Identified monetary policy shocks reduce household inflation uncertainty across categories. The transmission operates through prices: contractionary policy compresses realised price-change dispersion primarily in flexible food and gasoline categories, and households facing tighter price distributions report less uncertainty. This relationship appears both in the impulse responses and in direct sector-month regressions.

A two-sector noisy-information New Keynesian model turns this evidence into a policy prescription. Positive trend inflation makes price-change dispersion, and through it households' signal noise and posterior uncertainty, respond to policy at first order. The welfare-maximizing rule inverts the standard prescription across moments: it looks through the food-and-gas price level and targets core inflation, but leans against food-and-gas uncertainty, and does the reverse for services, because policy compresses dispersion only where prices adjust. Quantitatively, the rule moves the annualized policy rate about twice as strongly for a one-standard-deviation rise in food-and-gas uncertainty (roughly 100 basis points) as for core inflation. The stakes are not small: dropping the uncertainty response costs about 25 basis points of consumption each period, and the conventional headline rule about 35, in each case several times the classic cost of business cycles.

The measures open several avenues for future research. The rotating panel structure of the SCE is well suited to testing models of informational rigidities and belief formation; in Schnattinger and Srivastava (2025) we use the measures as an instrument for the pass-through of inflation uncertainty to wage-growth expectations. The sector-specific indices can also inform central bank communication and the design of rules that account for heterogeneous uncertainty across consumption categories.

References

Anand, R., E. S. Prasad, and B. Zhang (2015). What measure of inflation should a developing country central bank target? Journal of Monetary Economics 74, 102–116.

Andrade, P., E. Gautier, and E. Mengus (2023). What matters in households inflation expectations? Journal of Monetary Economics 138(C), 50–68.

Andreasen, M. M., J. Fernández-Villaverde, and J. F. Rubio-Ramírez (2018). The pruned state-space system for non-linear DSGE models: Theory and empirical applications. Review of Economic Studies 85(1), 1–49.

Anesti, N., V. Esady, and M. Naylor (2024). Food prices matter most: Sensitive household inflation expectations. Technical report, London School of Economics and Political Science.

Angeletos, G.-M. and Z. Huo (2021). Myopia and anchoring. American Economic Review 111(4), 1166–1200.

Angeletos, G.-M. and C. Lian (2018). Forward guidance without common knowledge. American Economic Review 108(9), 2477–2512.

Aoki, K. (2001). Optimal monetary policy responses to relative-price changes. Journal of Monetary Economics 48(1), 55–80.

Armantier, O., G. Topa, W. Van der Klaauw, and B. Zafar (2017). An overview of the survey of consumer expectations. Economic Policy Review 23(2), 51–72.

Ascari, G. (2004). Staggered prices and trend inflation: some nuisances. Review of Economic Dynamics 7(3), 642–667.

Ascari, G., S. Fasani, J. Grazzini, and L. Rossi (2023). Endogenous uncertainty and the macroeconomic impact of shocks to inflation expectations. Journal of Monetary Economics 140, S48–S63. Inflation: Drivers and Dynamics 2022.

Ascari, G. and A. M. Sbordone (2014). The macroeconomics of trend inflation. Journal of Economic Literature 52(3), 679–739.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. The quarterly journal of economics 131(4), 1593–1636.

Ball, L. (1992). Why does high inflation raise inflation uncertainty? Journal of Monetary Economics 29(3), 371–388.

Bauer, M. D. and E. T. Swanson (2023). A reassessment of monetary policy surprises and high-frequency identification. NBER Macroeconomics Annual 37(1), 87–155.

Ben-David, I., E. Fermand, C. M. Kuhnen, and G. Li (2018). Expectations uncertainty and household economic behavior. Technical report, National Bureau of Economic Research.

Benigno, P. (2004). Optimal monetary policy in a currency area. Journal of International Economics 63(2), 293–320.

Binder, C. C. (2017a). Economic policy uncertainty and household inflation uncertainty. The BE Journal of Macroeconomics 17(2), 20160048.

Binder, C. C. (2017b). Measuring uncertainty based on rounding: New method and application to inflation expectations. Journal of Monetary Economics 90, 1–12.

Binder, C. C. (2018). Inflation expectations and the price at the pump. Journal of Macroeconomics 58, 1–18.

Binetti, A., F. Nuzzi, and S. Stantcheva (2024). Peoples understanding of inflation. Technical report, Harvard University Working Paper.

Bloom, N. (2014). Fluctuations in uncertainty. Journal of economic Perspectives 28(2), 153–176.

Bodenstein, M., C. J. Erceg, and L. Guerrieri (2008). Optimal monetary policy with distinct core and headline inflation rates. Journal of Monetary Economics 55, S18–S33. Supplement.

Bullard, J. (2011). Measuring inflation: the core is rotten. Federal Reserve Bank of St. Louis Review 93(4), 223–234.

Caballero, R. J. and E. M. Engel (2007). Price stickiness in ss models: New interpretations of old results. Journal of Monetary Economics 54, 100–121.

Campos, C., M. McMain, and M. Pedemonte (2022). Understanding which prices affect inflation expectations. Economic Commentary 2022(06), 1–7.

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

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113(1), 1–45.

Clements, M. P. (2021). Rounding behaviour of professional macro-forecasters. International Journal of Forecasting 37(4), 1614–1631.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. Van Rooij (2023). How does consumption respond to news about inflation? field evidence from a randomized control trial. American Economic Journal: Macroeconomics 15(3), 109–152.

Coibion, O. and Y. Gorodnichenko (2012). What can survey forecasts tell us about information rigidities? Journal of Political Economy 120(1), 116–159.

Coibion, O. and Y. Gorodnichenko (2015a). Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review 105(8), 2644–2678.

Coibion, O. and Y. Gorodnichenko (2015b). 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.

Crump, R. K., S. Eusepi, A. Tambalotti, and G. Topa (2022). Subjective intertemporal substitution. Journal of Monetary Economics 126, 118–133.

Cukierman, A. and A. H. Meltzer (1986). A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. Econometrica 54(5), 1099–1128.

Dietrich, A. M. (2024). Consumption categories, household attention, and inflation expectations: Implications for optimal monetary policy. Journal of monetary economics 147, 103594.

Dominitz, J. and C. F. Manski (1997). Using expectations data to study subjective income expectations. Journal of the American statistical Association 92(439), 855–867.

Dovern, J. and A. Glas (2025). Uncertainty and rounding in expectation surveys. Discussion Paper 22/2025, Deutsche Bundesbank.

DAcunto, F., E. Charalambakis, D. Georgarakos, G. Kenny, J. Meyer, and M. Weber (2024, May). Household inflation expectations: An overview of recent insights for monetary policy. Technical Report 32488, National Bureau of Economic Research. NBER Working Paper (W32488).

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

Engelberg, J., C. F. Manski, and J. Williams (2009). Comparing the point predictions and subjective probability distributions of professional forecasters. Journal of Business & Economic Statistics 27(1), 30–41.

Erceg, C. J. and A. T. Levin (2006). Optimal monetary policy with durable consumption goods. Journal of Monetary Economics 53(7), 1341–1359.

Eusepi, S., B. Hobijn, and A. Tambalotti (2011). Condi: A cost-of-nominal-distortions index. American Economic Journal: Macroeconomics 3(3), 53–91.

Fama, E. F. and J. D. MacBeth (1973). Risk, return, and equilibrium: Empirical tests. Journal of political economy 81(3), 607–636.

Fischer, J. J., C. Herler, and P. Schnattinger (2025). When the fog clears: The effect of reduced inflation uncertainty on households' financial behaviour. Staff Working Paper 1133, Bank of England.

Flynn, J. P., G. Nikolakoudis, and K. A. Sastry (2026). Supply function choice and aggregate supply. American Economic Review. Forthcoming.

Friedman, M. (1977). Nobel lecture: inflation and unemployment. Journal of political economy 85(3), 451–472.

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2022). Expectations with endogenous information acquisition: An experimental investigation. Review of Economics and Statistics 104(5), 1059–1078.

Galí, J. (2015). Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications (2nd ed.). Princeton University Press.

Georgarakos, D., Y. Gorodnichenko, O. Coibion, and G. Kenny (2024). The causal effects of inflation uncertainty on households' beliefs and actions. Working Paper 33014, National Bureau of Economic Research.

Giustinelli, P., C. F. Manski, and F. Molinari (2022). Tail and center rounding of probabilistic expectations in the health and retirement study. Journal of Econometrics 231(1), 265–281. Annals Issue: Subjective Expectations &ăProbabilities in Economics.

Golosov, M. and J. Lucas, Robert E. (2007). Menu costs and phillips curves. Journal of Political Economy 115(2), 171–199.

Hoffmann, M., E. Moench, L. Pavlova, and G. Schultefrankenfeld (2022). Would households understand average inflation targeting? Journal of Monetary Economics 129, S52–S66.

Hughes, J. E., C. R. Knittel, and D. Sperling (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. The Energy Journal 29(1), 113–134.

Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprisesthe role of information shocks. American Economic Journal: Macroeconomics 12(2), 1–43.

Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring uncertainty. American Economic Review 105(3), 1177–1216.

Kim, G. and C. Binder (2023). Learning-through-survey in inflation expectations. American Economic Journal: Macroeconomics 15(2), 254–278.

Kostyshyna, O. and L. Petersen (2024). The effect of inflation uncertainty on household expectations and spending. Technical report, National Bureau of Economic Research.

Krifka, M. (2007). Approximate interpretation of number words. Humboldt-Universität zu Berlin, Philosophische Fakultät II.

Lamla, M. J. and D. V. Vinogradov (2019). Central bank announcements: Big news for little people? Journal of Monetary Economics 108, 21–38.

Lauper, C. and G. Mangiante (2025). Monetary policy shocks and inflation inequality. International Journal of Central Banking 21(4).

Lorenzoni, G. (2009). A theory of demand shocks. American Economic Review 99(5), 2050–2084.

Lucas, Robert E., J. (2003). Macroeconomic priorities. American Economic Review 93(1), 1–14.

Maćkowiak, B., F. Matějka, and M. Wiederholt (2023). Rational inattention: A review. Journal of Economic Literature 61(1), 226–273.

Maćkowiak, B. and M. Wiederholt (2009). Optimal sticky prices under rational inattention. American Economic Review 99(3), 769–803.

Maćkowiak, B. and M. Wiederholt (2015). Business cycle dynamics under rational inattention. The Review of Economic Studies 82(4), 1502–1532.

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

Mankiw, N. G. and R. Reis (2002). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. The Quarterly Journal of Economics 117(4), 1295–1328.

Mankiw, N. G. and R. Reis (2003). What measure of inflation should a central bank target? Journal of the European Economic Association 1(5), 1058–1086.

Melosi, L. (2017). Signalling effects of monetary policy. Review of Economic Studies 84(2), 853–884.

Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. Quarterly Journal of Economics 123(4), 1415–1464.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018). The elusive costs of inflation: Price dispersion during the U.S. great inflation. Quarterly Journal of Economics 133(4), 1933–1980.

Nimark, K. (2008). Dynamic pricing and imperfect common knowledge. Journal of Monetary Economics 55(2), 365–382.

Nord, L. (2022). Who cares about inflation? endogenous expectation formation of heterogeneous households. Manuscript titled "Endogenous Expectation Formation of Heterogeneous Households" (July 14, 2022).

Oh, J. (2020). The propagation of uncertainty shocks: Rotemberg versus Calvo. International Economic Review 61(3), 1097–1113.

Paciello, L. and M. Wiederholt (2014). Exogenous information, endogenous information, and optimal monetary policy. Review of Economic Studies 81(1), 356–388.

Reiche, L. and A. Meyler (2022). Making sense of consumer inflation expectations: The role of uncertainty. Working Paper 2733, European Central Bank. ECB Working Paper.

Schnattinger, P. and P. Srivastava (2025). Household inflation uncertainty and wagegrowth expectations. Working paper, Bank of England. Unpublished working paper.

Shapley, L. S. (1953). A value for n-person games. In H. W. Kuhn and A. W. Tucker (Eds.), Contributions to the Theory of Games II, Volume 28 of Annals of Mathematics Studies, pp. 307–317. Princeton University Press.

Sheremirov, V. (2020). Price dispersion and inflation: New facts and theoretical implications. Journal of Monetary Economics 114, 59–70.

Shorrocks, A. F. et al. (2013). Decomposition procedures for distributional analysis: a unified framework based on the shapley value. Journal of Economic Inequality 11(1), 99–126.

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

Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. American Economic Review 97(3), 586–606.

Stantcheva, S. (2024). Why do we dislike inflation? Technical report, National Bureau of Economic Research.

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

Zhao, Y. (2023). Internal consistency of household inflation expectations: Point forecasts vs. density forecasts. International Journal of Forecasting 39(4), 1713–1735.

Zhong, C. and J. Huo (2024). Does age matter in inflation uncertainty? Economics Letters 242, 111834.

Online Appendix for "Sector-Specific Household Inflation Uncertainty and Monetary Policy"

Philip Schnattinger Prachi Srivastava

A Derivations for the Framework of Section 2

A.1 Item-level Bayesian updating

With the prior $\pi_{j,t+1} \sim \mathcal{N}(\mu_{i,j}, \sigma^2_{i,j})$ and the signal (1), conjugate normal updating gives the posterior precision (2) and the posterior mean

$$\mathbb{E}_i(\pi_{j,t+1} \mid s_{i,j,t}) = \mathrm{Var}_i(\pi_{j,t+1} \mid s_{i,j,t}) \left(\frac{\mu_{i,j}}{\sigma^2_{i,j}} + \frac{s_{i,j,t}}{\sigma^2_{\eta,i,j,t}} \right), \tag{A.1}$$

the precision-weighted average of prior and signal.

B SCE Question Wording

This appendix records the SCE question wording used to construct the aggregate and sector-specific belief measures. Question Q8 elicits the respondent's point forecast for aggregate inflation, Q9 elicits the subjective density forecast used to construct density-based aggregate uncertainty, and C4 elicits the sector-specific price forecasts used to construct the RNRI measures.

Inflation point forecast (Q8v2part2).

"What do you expect the rate of [inflation/deflation] to be over the next 12 months? Please give your best guess."

Over the next 12 months, I expect the rate of [inflation/deflation] to be ...

Inflation density forecast (Q9).

Now we would like you to think about the different things that may happen to inflation over the next 12 months.

What would you say is the percent chance that, over the next 12 months ...

Bin	Chance	Bin	Chance
Go up by 12% or more	– %	Go down by 0% to 2%	– %
Go up by 8% to 12%	– %	Go down by 2% to 4%	– %
Go up by 4% to 8%	– %	Go down by 4% to 8%	– %
Go up by 2% to 4%	– %	Go down by 8% to 12%	– %
Go up by 0% to 2%	– %	Go down by 12% or more	– %

Sector-specific price forecasts (C4info).

Twelve months from now, what do you think will have happened to the price of the following items?

I expect ...

Item	Increase (%)		Decrease (%)
The price of a gallon of gas	– %	OR	– %
The price of food	– %	OR	– %
The price of medical care	– %	OR	– %
The cost of a college education	– %	OR	– %
The cost of renting a typical house/apartment	– %	OR	– %

C Estimating Uncertainty: RNRI

This appendix states the mixture model underlying the RNRI uncertainty proxy χ_{jit} of Equation (4).

Let R_{ijt} denote the reported forecast of respondent i for sector j in month t. The RNRI model treats the reported value as a rounded version of an underlying latent forecast, and allows respondents to differ in the coarseness with which they report their beliefs: low-RNRI types report on a fine grid, high-RNRI types on a coarser grid of multiples of five.

A low-RNRI type draws a latent forecast from $\mathcal{N}\left(\mu_{jt}^{l}, (\sigma_{jt}^{l})^2\right)$ and reports it to the nearest integer, so the probability of observing report r is

$$\psi_{jt}^{l}(r) \;=\; \Pr(R_{ijt} = r \mid l) \;=\; \int_{r-0.5}^{r+0.5} \frac{1}{\sigma_{jt}^{l}\sqrt{2\pi}} \exp\left\{-\frac{(x-\mu_{jt}^{l})^2}{2(\sigma_{jt}^{l})^2}\right\} dx, \qquad r = \ldots, -1, 0, 1, \ldots \tag{C.2}$$

A high-RNRI type draws from $\mathcal{N}\left(\mu_{jt}^{h}, (\sigma_{jt}^{h})^2\right)$ and reports to the nearest multiple of five, so

$$\psi_{jt}^{h}(r) \;=\; \Pr(R_{ijt} = r \mid h) \;=\; \int_{r-2.5}^{r+2.5} \frac{1}{\sigma_{jt}^{h}\sqrt{2\pi}} \exp\left\{-\frac{(x-\mu_{jt}^{h})^2}{2(\sigma_{jt}^{h})^2}\right\} dx, \qquad r = \ldots, -5, 0, 5, \ldots \tag{C.3}$$

The two mass functions differ only in the reporting grid: low-RNRI types place probability mass on one-unit intervals around each integer, high-RNRI types on five-unit intervals around each multiple of five. In both cases the intervals tile the real line, so ψ_{jt}^{l} and ψ_{jt}^{h} are proper probability mass functions on their respective supports. The model therefore interprets coarse reporting as informative about the respondent's latent uncertainty type, following the idea that rounding is more likely when beliefs are less precise (Giustinelli et al., 2022).

The differing supports carry a sharp implication. Since $\psi_{jt}^{h}(r) = 0$ whenever r is not a multiple of five, a non-round response can only have come from the low-RNRI component: such a respondent is classified as low-uncertainty with certainty, and $\chi_{jit} = 0$. The proxy varies strictly between zero and one only across round responses, graded there by the

relative likelihood of the report under the two components. This is what separates χ_{jit} from the raw M5 indicator, which treats every round response alike.

Responses in sector j and month t are generated by the two-component mixture

$$\psi_{jt}(r \mid \theta_{jt}) \;=\; \lambda_{jt}\,\psi^{h}_{jt}(r) \;+\; \big(1-\lambda_{jt}\big)\,\psi^{l}_{jt}(r), \qquad \theta_{jt} \equiv \{\lambda_{jt},\, \mu^{l}_{jt},\, \mu^{h}_{jt},\, \sigma^{l}_{jt},\, \sigma^{h}_{jt}\}, \tag{C.4}$$

where $\lambda_{jt} = N^{h}_{jt}/N_{jt}$ is the share of high-RNRI types among the $N_{jt} = N^{h}_{jt} + N^{l}_{jt}$ respondents. Given $\{R_{ijt}\}_{i=1}^{N_{jt}}$, we estimate θ_{jt} separately for each sector-month by maximum likelihood,

$$\hat{\theta}_{jt} \;=\; \arg\max_{\theta_{jt}} \sum_{i=1}^{N_{jt}} \log \psi_{jt}(R_{ijt} \mid \theta_{jt})\,. \tag{C.5}$$

Substituting $\hat{\theta}_{jt}$ into Bayes' rule delivers the posterior probability χ_{jit} of Equation (4) that a given response comes from the high-RNRI component.

D Replication of Uncertainty Indices

As a validation exercise, we first replicate the Michigan Consumer Survey (MCS) inflation-uncertainty index of Binder (2017b). We estimate the index by maximum likelihood in R using the L-BFGS algorithm and compare our replicated series with the published MCS index.[17] Figure D.1 shows that the replicated series closely matches the published index; the correlation between the two series is 0.98.

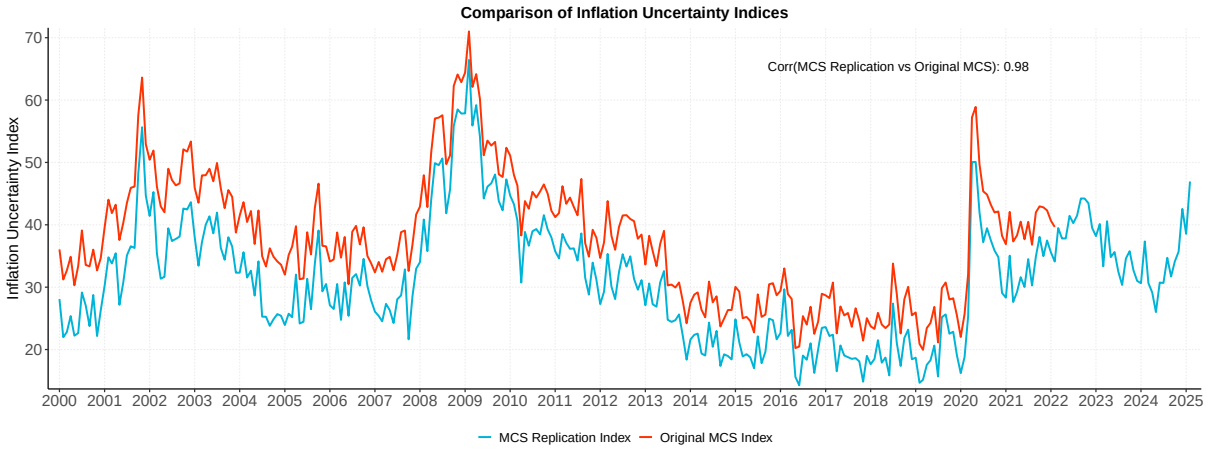


Figure D.1: Replication of the MCS Inflation-Uncertainty Index

Note: This figure compares our replicated MCS inflation-uncertainty index with the index published by Binder (2017b). The published data are from Inflation Uncertainty Index.

We then apply the same maximum-likelihood procedure to the Survey of Consumer Expectations (SCE) point forecasts from question Q8v2part2 to construct a parallel SCE-based inflation-uncertainty index. Figure D.2 compares the SCE-based index with the MCS index over their common sample. The two series co-move strongly: both capture the sharp increase in inflation uncertainty during the pandemic and its subsequent decline.

[17] Data available at Inflation Uncertainty Index.

Over June 2013–June 2024, their correlation is 0.81. The remaining differences likely reflect differences in survey wording, sample composition, timing, and implementation.

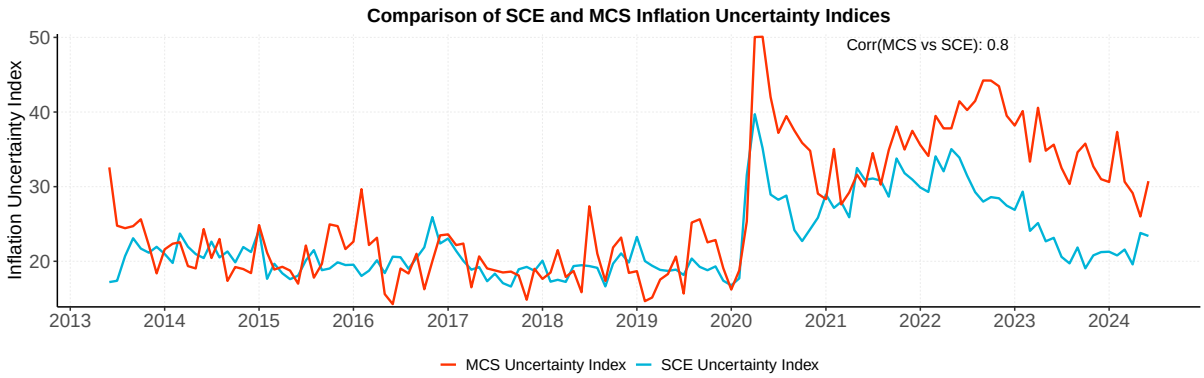


Figure D.2: SCE-Based and MCS Inflation-Uncertainty Indices

Note: This figure compares the SCE-based inflation-uncertainty index, constructed from Q8v2part2 point forecasts using the same maximum-likelihood procedure, with the MCS index of Binder (2017b). The published MCS data are from Inflation Uncertainty Index.

E Estimated Sectoral Distribution Parameters

This section reports the fitted sector-level parameters behind the RNRI uncertainty indices used in the main text. Section 3.3 constructs the main uncertainty proxy, χ_{jit}, as the posterior probability that respondent i belongs to the high-RNRI type in sector j and month t. The objects plotted here are the maximum-likelihood estimates of the two type-specific forecast distributions that enter that posterior probability.

For each sector and month, we estimate the low- and high-RNRI type-specific means and standard deviations. Figure E.3 plots the estimated means, μ_{jlt} and μ_{jht}, while Figure E.4 plots the estimated standard deviations, σ_{jlt} and σ_{jht}. The solid lines are the original-sample maximum-likelihood estimates. The shaded bands summarize sampling uncertainty from the bootstrap: in each bootstrap replication the mixture model is re-estimated, and the 95% band is formed from the empirical distribution of the resulting parameter estimates. Thus, the figures should be read as diagnostics for the fitted low- and high-RNRI distributions, while χ_{jit} remains the individual-level uncertainty index used in the empirical analysis. Across sectors, the high-RNRI type generally has both higher mean expected price changes and larger dispersion than the low-RNRI type.

F Robustness using different measures from the SCE

The aggregate comparison is informative, but sectoral-level uncertainty is the object of interest in this paper, and the SCE does not elicit sector-specific density forecasts. We therefore use the aggregate density forecast as an external benchmark and exploit within-respondent variation in sector-level RNRI. Specifically, we estimate

$$\sigma_{it}^{\pi,\text{dens}} = \alpha_i + \delta_t + \sum_{j \in \mathcal{J}} \beta_j \chi_{jit} + \varepsilon_{it}, \tag{F.6}$$

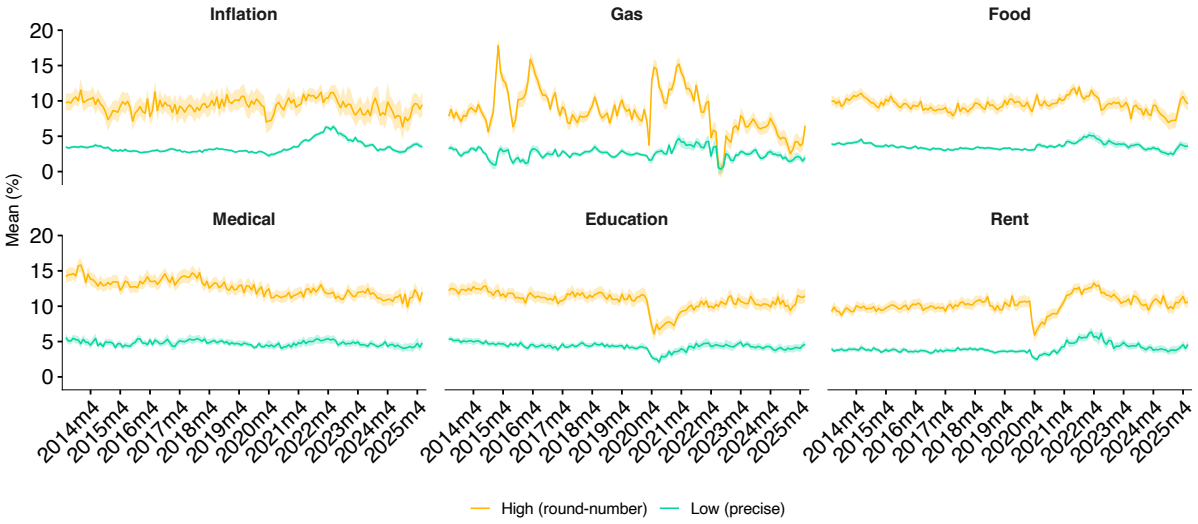


Figure E.3: Estimated Sectoral Means for High- and Low-RNRI Types

Note: The figure shows maximum-likelihood estimates of the sector-month means, μ_{jlt} and μ_{jht}, for low- and high-RNRI types. The plotted lines are the original-sample estimates; shaded bands report bootstrapped 95% confidence intervals based on re-estimating the mixture model across bootstrap samples. The five sectors are gasoline, food, medical care, education, and rent. The sample runs from June 2013 to June 2025.

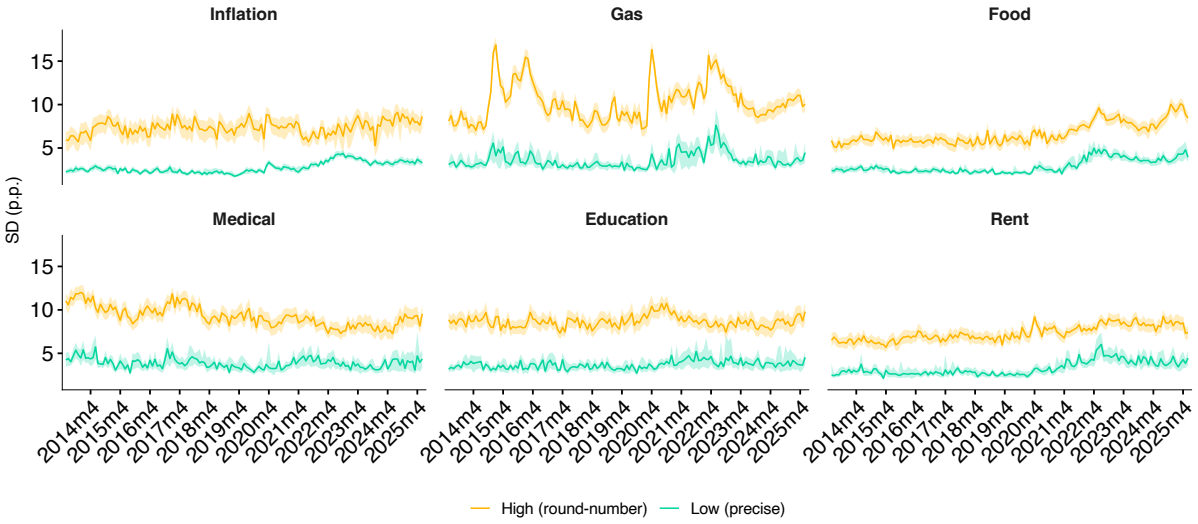


Figure E.4: Estimated Sectoral Standard Deviations for High- and Low-RNRI Types

Note: The figure shows maximum-likelihood estimates of the sector-month standard deviations, σ_{jlt} and σ_{jht}, for low- and high-RNRI types. The plotted lines are the original-sample estimates; shaded bands report bootstrapped 95% confidence intervals based on re-estimating the mixture model across bootstrap samples. The five sectors are gasoline, food, medical care, education, and rent. The sample runs from June 2013 to June 2025.

where α_i are respondent fixed effects and δ_t are survey-month fixed effects. This specification absorbs stable differences in numeracy, education, attention, and reporting style, as well as aggregate shocks common to all respondents in a month. The coefficients are therefore identified from changes in the same respondent's sectoral-level imprecision

over time. Appendix Table F.1 shows that all five sectors χ measures positively and significantly predict density-based aggregate uncertainty under this within-respondent specification, with the largest coefficients for food and rent. This pattern is difficult to reconcile with an interpretation in which RNRI only captures fixed rounding habits.

F.1 Using user fixed effects

In Table F.1 we re-estimate the baseline dummy and proxy specifications, this time including individual-level fixed effects in place of the demographic and financial-condition controls. By exploiting only within-individual variation, this approach controls for time-invariant omitted variables. Consistent with our baseline results, columns (1) and (2) show that the continuous-probability uncertainty measure yields a higher R^2 than the binary dummy, with food and rent uncertainty exhibiting the largest coefficients. Likewise, in columns (3) and (4), where we use the standard deviation from the subjective probability distribution from SCE as the measure of uncertainty, we again find the coefficient on gas uncertainty is similar in magnitude to before, while food and rent uncertainty are the strongest predictors of individual inflation-uncertainty perceptions.

F.2 Forecast-Error Validation: Robustness to Forecast-Level Controls

A mechanical concern with the forecast-error validation of Table 2 is that high-χ respondents report larger point forecasts, so their absolute errors could be larger in a low-inflation sample even absent any link between rounding and subjective uncertainty. Table F.2 re-estimates the specification adding the absolute deviation of the respondent's point forecast from the month-category median (column 2) or the forecast level itself (column 3). The level channel accounts for roughly half of the raw coefficient, and the remainder stays large and precise in every category.

Table F.1: Within-Respondent Validation of Sectoral RNRI

	Inflation expectation		Density-based uncertainty	
	M5 dummy	RNRI posterior	M5 dummy	RNRI posterior
	(1)	(2)	(3)	(4)
Sectoral rounding indicators				
$M5^{gas}$	0.22***		0.11***	
	(0.03)		(0.01)	
$M5^{food}$	0.47***		0.13***	
	(0.03)		(0.01)	
$M5^{med}$	0.10***		0.05***	
	(0.03)		(0.01)	
$M5^{edu}$	0.14***		0.08***	
	(0.03)		(0.01)	
$M5^{rent}$	0.24***		0.10***	
	(0.03)		(0.01)	
Sectoral RNRI posterior probabilities				
χ^{gas}		0.20***		0.11***
		(0.04)		(0.01)
χ^{food}		0.89***		0.24***
		(0.04)		(0.02)
χ^{med}		0.14***		0.08***
		(0.04)		(0.01)
χ^{edu}		0.18***		0.11***
		(0.04)		(0.01)
χ^{rent}		0.36***		0.15***
		(0.04)		(0.01)
Respondent fixed effects	Yes	Yes	Yes	Yes
Survey-month fixed effects	Yes	Yes	Yes	Yes
Observations	145,045	145,045	145,045	145,045
R^2	0.621	0.622	0.741	0.742

Notes: Each cell reports the coefficient with its standard error in parentheses below. Columns (1) and (3) use indicators for multiple-of-five forecasts; columns (2) and (4) use the RNRI posterior probability χ. All specifications include respondent and survey-month fixed effects. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F.2: Forecast-Error Validation: Robustness to Forecast-Level Controls

Category	Baseline	+ |dev. from month median|	+ forecast level	N
Gas	2.465***	0.979***	1.022***	127,649
	(0.079)	(0.076)	(0.079)	
Food	5.105***	2.537***	2.173***	126,926
	(0.051)	(0.036)	(0.064)	
Medical	6.306***	4.283***	1.033***	126,926
	(0.061)	(0.045)	(0.050)	
Education	5.808***	2.773***	2.087***	126,926
	(0.053)	(0.027)	(0.055)	
Rent	4.285***	2.088***	1.331***	126,926
	(0.048)	(0.018)	(0.059)	

Notes: Coefficient on the rounding-based uncertainty proxy χ_{jit} in regressions of the absolute realized forecast error on χ_{jit} with respondent and month fixed effects (the specification of Table 2), adding the absolute deviation of the respondent's point forecast from the month-category median (column 2) or the forecast level itself (column 3). Standard errors clustered by respondent in parentheses. Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

G Monetary Policy Surprise Series for Local Projections

G.1 Jarociński–Karadi Shock Series

This appendix documents the monetary policy surprise series used in the local projections. The baseline specification uses the `MP_pm` series of Jarociński and Karadi (2020), which separates monetary policy shocks from central-bank information shocks using sign restrictions. The underlying policy surprise is the first principal component of high-frequency changes in interest-rate derivatives with maturities from one month to one year around FOMC announcements.

Figure G.5 plots the monthly shock series over the estimation period. Positive values correspond to contractionary monetary policy surprises and negative values to expansionary surprises.

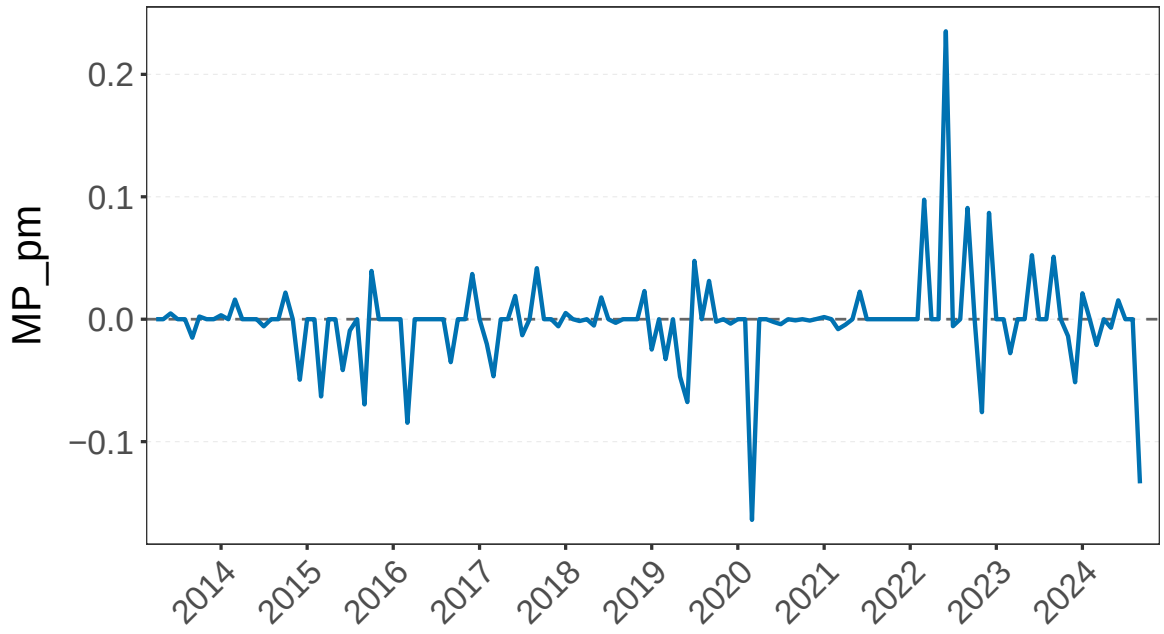


Figure G.5: Jarociński–Karadi Monetary Policy Surprise Series

Note: The figure plots the monthly `MP_pm` monetary policy shock series from Jarociński and Karadi (2020). The series is constructed from high-frequency financial-market surprises around FOMC announcements and is scaled in percentage points. The series runs from June 2013 to January 2024.

G.2 Bauer–Swanson Orthogonalized Surprise Series

As a robustness check, we replace the baseline series with the orthogonalized high-frequency monetary policy surprises of Bauer and Swanson (2023). Their raw surprise is the first principal component of 30-minute changes in money-market futures rates around FOMC announcements and Fed Chair speeches, capturing revisions to the expected policy-rate path over the next year. Since raw high-frequency surprises partly reflect the Federal Reserve's systematic response to public news, Bauer and Swanson (2023) orthogonalize the series with respect to pre-announcement macroeconomic and

financial conditions. We use the residual series, $\texttt{MPS_ORTH}_t$, as a cleaner measure of unanticipated policy surprises.

We aggregate the event-level series to monthly frequency and re-estimate the symmetric local-projection specification in Equation (14). The figures below therefore isolate sensitivity to the shock measure: the regression specification is held fixed while the monetary policy surprise series changes.

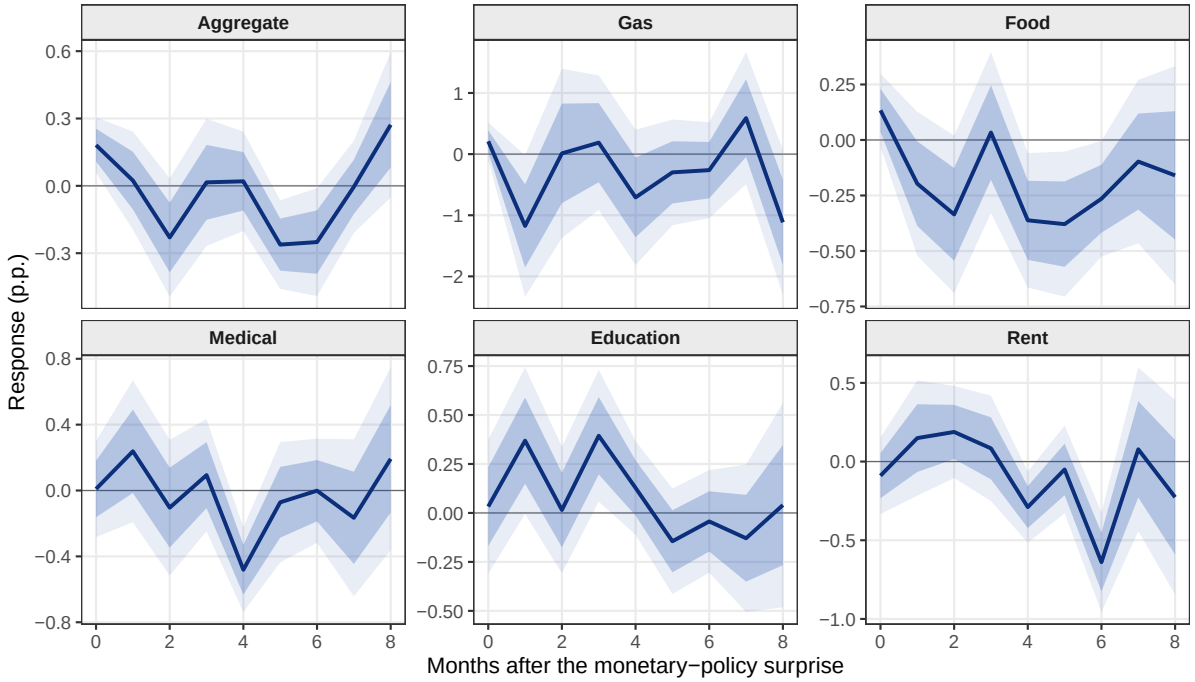


Figure G.6: Bauer–Swanson Monetary Policy Surprises and Sectoral Price Expectations

Note: Estimated impulse responses of household aggregate and sectoral price expectations to a 25 basis point monetary policy surprise, identified using the orthogonalized Bauer and Swanson (2023) series. The specification matches the baseline symmetric local projections: an 8-month horizon, the controls in Equation (14), respondent fixed effects, the outcome-date COVID fixed effects (March–June 2020), month-of-year seasonal fixed effects, and standard errors clustered by survey month. The sample covers June 2013 to December 2023, with the end date determined by the availability of the surprise series. Shaded bands represent the 66.7% and 90% confidence intervals.

G.3 Dispersion and Household Uncertainty: Direct Evidence

The local projections of Section 5 establish that a monetary shock moves both realized sectoral price dispersion and household sectoral uncertainty. This subsection estimates the middle link directly. In a sector-month panel of the five SCE categories, we regress household uncertainty on the sector's realized BLS price dispersion,

$$y_{j,t} \;=\; \alpha_j \;+\; \gamma_t \;+\; \beta\, \mathrm{Disp}_{j,t}^{BLS} \;+\; u_{j,t},$$

where $y_{j,t}$ is the monthly mean (or cross-sectional dispersion) of household category-j uncertainty, $\mathrm{Disp}_{j,t}^{BLS}$ is the interquartile range of year-on-year inflation across the detailed BLS items in category j, and α_j and γ_t are sector and month fixed effects, so that β is identified from within-month, within-sector variation. Table G.3 reports the pooled estimates: months in which a sector's price dispersion is wider are months in which

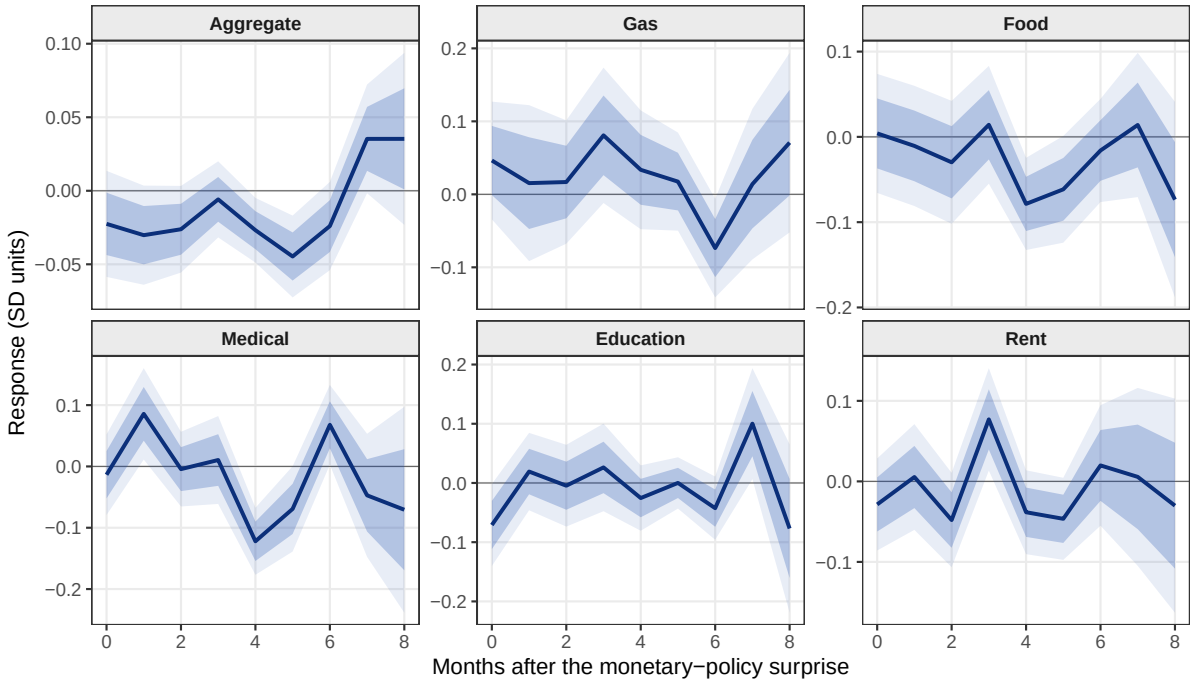


Figure G.7: Bauer–Swanson Monetary Policy Surprises and Inflation Uncertainty

Note: Estimated impulse responses of household aggregate and sectoral inflation uncertainty, standardized to have mean zero and unit variance, to a 25 basis point monetary policy surprise identified using the orthogonalized Bauer and Swanson (2023) series. The specification matches the baseline symmetric local projections: an 8-month horizon, the controls in Equation (14), respondent fixed effects, the outcome-date COVID fixed effects (March–June 2020), month-of-year seasonal fixed effects, and standard errors clustered by survey month. The responses are similar to the baseline in sign and sectoral ordering. The sample covers June 2013 to December 2023, with the end date determined by the availability of the surprise series. Shaded bands represent the 66.7% and 90% confidence intervals.

households are both more uncertain about that sector on average and more dispersed in their uncertainty, and disagreement about expected sectoral inflation is wider. Table G.4 estimates the uncertainty-dispersion regression sector by sector: the association is significant for food, rent, and medical care, and absent for gas and education, where category uncertainty is dominated by a single common price that households observe directly. This is the middle link the signal-extraction framework of Section 2 formalizes.

G.4 Asymmetry of the Uncertainty Response

The baseline local projections pool tightening and easing shocks, because the model counterpart of the estimates is a perturbation impulse response, which is odd in the shock by construction: the symmetric coefficient is the model-consistent average effect. The data, however, do not impose symmetry. Figure G.8 splits the shock into tightening and easing components (estimated under the baseline specification, including the outcome-date COVID fixed effects and month-of-year seasonal fixed effects), and Table G.5 reports the symmetry test: the coefficient δ_h on the tightening component in a regression that includes both the pooled shock and the tightening part, so that δ_h measures the departure of the tightening response from the symmetric benchmark. Symmetry is rejected at the trough horizons for gas and food: monetary tightenings lower flexible-sector uncertainty by roughly twice the pooled average, while the easing side is weaker and less precisely

Table G.3: Household Uncertainty and BLS Price Dispersion: Pooled Sector-Month Panel

	Dependent Variable		
Regressor	**Expectation Disagreement** (1)	**RNRI Mean Uncertainty** (2)	**RNRI Uncertainty Disagreement** (3)
Price dispersion	0.5771***	0.0115***	0.0185***
	(0.0808)	(0.0015)	(0.0069)
Sector fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
Observations	725	725	725
R^2	0.4495	0.8717	0.3588

Note: The sample is a sector-month panel covering five SCE categories at a monthly frequency. The dependent variables are the cross-sectional interquartile range of one-year-ahead sectoral price expectations (expectation disagreement), the monthly mean of household sectoral uncertainty (ζ), and the cross-sectional interquartile range of household sectoral uncertainty (uncertainty disagreement). The regressor is the interquartile range of year-on-year inflation across the detailed BLS items within each category. All specifications include sector and month fixed effects. Standard errors clustered by month are reported in parentheses. Significance levels are $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table G.4: Household Uncertainty Dispersion and BLS Price Dispersion by Sector

	Household Uncertainty Dispersion				
Regressor	**Gas** (1)	**Food** (2)	**Rent** (3)	**Medical** (4)	**Education** (5)
BLS price dispersion	-0.0132	0.0593***	0.4945**	0.0514***	0.0010
	(0.0122)	(0.0085)	(0.2245)	(0.0171)	(0.0027)
Constant	0.9492***	0.5043***	0.7635***	0.7476***	0.8921***
	(0.0225)	(0.0409)	(0.0587)	(0.0550)	(0.0110)
Observations	145	145	145	145	145
R^2	0.0139	0.1668	0.0320	0.0483	0.0003

Note: Each column reports a separate sector-level time-series regression of the monthly cross-sectional dispersion in household uncertainty for category j on the corresponding category's BLS item-level price dispersion. Household uncertainty dispersion is measured by the interquartile range of respondent-level category uncertainty. BLS price dispersion is measured by the interquartile range of year-on-year inflation across detailed BLS items within the category. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels are $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

estimated. The pooled coefficients used in the main text are therefore conservative for the effect of a tightening.

H Plausibility Check: Inflation Scarring and Price Exposure

As a plausibility check on the RNRI measures, we ask whether groups with different exposure to price volatility also display different levels of sector-specific uncertainty. The exercise is descriptive rather than causal. The motivation is the experience-based learning evidence of Malmendier and Nagel (2016), who show that inflation experiences leave

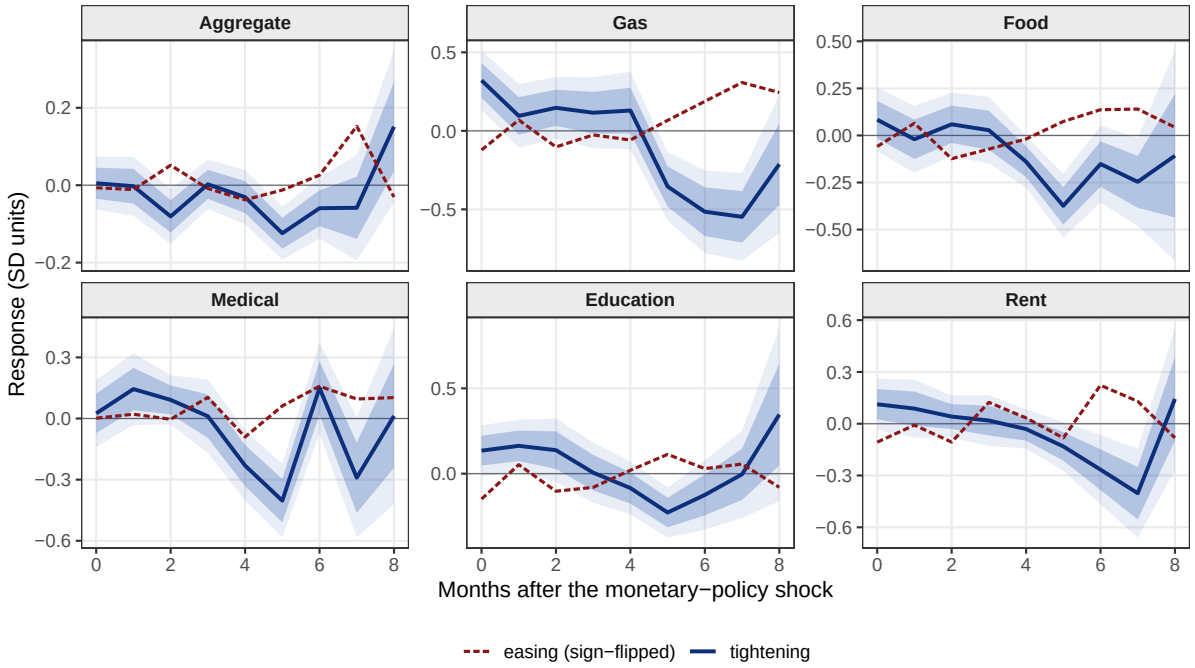


Figure G.8: Tightening versus easing: sectoral uncertainty responses

Note: Impulse responses of household inflation uncertainty (standardized) to a 25 basis point shock, estimated separately for tightening (solid, with the nested 66.7% and 90% bands) and easing (dashed, sign-flipped so that under symmetry the two lines coincide) components of the Jarociński and Karadi (2020) series. Specification and sample as in the baseline, including the outcome-date COVID fixed effects (March–June 2020) and month-of-year seasonal fixed effects.

Table G.5: Symmetry test: departure of the tightening response from the symmetric benchmark

Outcome	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$	$h=7$	$h=8$
Aggregate	$+0.01$	$+0.01$	-0.13^{**}	$+0.01$	$+0.01$	-0.11^{*}	-0.09	-0.21^{*}	$+0.18$
Gas	$+0.44^{**}$	$+0.03$	$+0.25$	$+0.14$	$+0.19$	-0.42^{*}	-0.70^{***}	-0.86^{***}	-0.46
Food	$+0.14$	-0.08	$+0.18$	$+0.10$	-0.12	-0.45^{***}	-0.29	-0.39^{*}	-0.15
Medical	$+0.02$	$+0.12$	$+0.10$	-0.09	-0.14	-0.46^{***}	-0.01	-0.39	-0.09
Rent	$+0.22$	$+0.10$	$+0.15$	-0.11	-0.06	-0.05	-0.49^{***}	-0.53^{**}	$+0.23$

Note: Each cell reports δ_h, the coefficient on the tightening component of the shock in a regression that also includes the pooled shock, scaled to 25 basis points; $\delta_h = 0$ under symmetry. Negative entries mean the tightening response is more negative than the symmetric benchmark. Specification and sample as in the baseline, including the outcome-date COVID fixed effects (March–June 2020) and month-of-year seasonal fixed effects; standard errors clustered by survey month. *, **, ***: significant at the 10, 5, and 1 percent level.

persistent scars on inflation expectations. If RNRI captures meaningful variation in subjective uncertainty, we would expect higher RNRI values among cohorts or demographic groups whose lived experience or consumption exposure makes particular prices more salient.

This check also relates to Zhong and Huo (2024), who use SCE data and show that RNRI- and density-based uncertainty measures can imply different age gradients. That distinction is useful for our purposes: density questions require respondents to allocate probabilities across bins, while RNRI is inferred from the coarseness of the point forecast. We therefore interpret the demographic patterns below as evidence that RNRI contains a behavioral signal about imprecision that is complementary to density-based measures.

We begin with age and gender. Figure H.9 shows a clear age gradient: older respondents report higher RNRI uncertainty across sectors. Gender differences are also sector-specific. Women report higher uncertainty for food, education, medical care, rent, and the aggregate RNRI-based measure. Gasoline is the main exception: men report higher gas-price uncertainty in every age cohort.

These patterns are consistent with differences in exposure and salience. For gasoline, driving intensity is a natural candidate: statistics from the National Household Travel Survey indicate that men drive more miles than women.[18] More frequent exposure to fuel purchases may make gasoline-price volatility more salient. By contrast, higher uncertainty among women for food and household-budget categories is consistent with the idea that everyday shopping and budgeting are important channels through which households form price perceptions. Because energy prices also pass through to other expenses, gas-price volatility may affect uncertainty about non-gas categories as well.

The age gradient is also consistent with differential exposure and vulnerability. Older respondents have lived through more episodes of high or volatile inflation and may place greater weight on essential, non-discretionary expenditures. Age profiles of energy, transportation, housing, and medical spending differ substantially over the life cycle,[19] so the same price shock can have different salience across cohorts. Older households may also be more exposed to rent and medical-cost risk, both of which can raise uncertainty about future household budgets.[20]

In the language of the signal-extraction framework in Section 2, the demographic gradients can be interpreted as heterogeneity in prior beliefs, $(\mu_{i,j}, \sigma^2_{i,j})$. Individuals with different consumption baskets and lived price histories enter the forecasting problem with different priors, consistent with the experience-based learning channel in Malmendier and Nagel (2016). The fact that RNRI varies systematically with these exposure margins supports its interpretation as a proxy for item-level uncertainty.

[18] U.S. Federal Highway Administration, National Household Travel Survey; the gender gap in annual vehicle-miles traveled is offered as illustrative context rather than direct evidence on uncertainty.

[19] Age profiles of expenditure shares are documented in the BLS Consumer Expenditure Survey; the link to perceived price uncertainty is interpretive rather than causal.

[20] Older Americans Increasingly Vulnerable to Rising Rents.

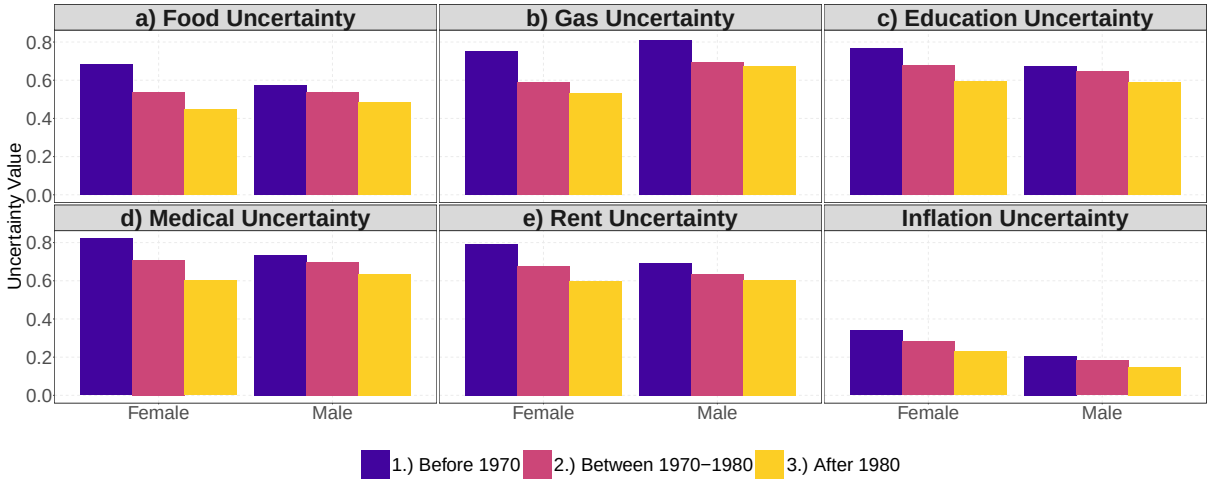


Figure H.9: Age, Gender, and Sector-Specific Inflation Uncertainty

Note: This figure shows average sector-specific uncertainty by age group and gender using the RNRI-based uncertainty measure. Sample period: June 2013 to June 2025.

Figure H.10 examines how sector-specific uncertainty changed after the onset of the COVID-19 pandemic. For each age–sector pair, the figure plots the proportional change in average uncertainty, (After/Before) − 1. Uncertainty rises for all age cohorts and all categories, but the increase is not uniform. The largest post-COVID increases for food, gasoline, rent, and aggregate RNRI uncertainty are concentrated among respondents born between 1970 and 1980. Respondents born before 1970 show the largest increases for education and medical care, while the youngest cohort, born after 1980, generally exhibits smaller increases, especially for rent.

These patterns suggest that the pandemic and subsequent inflation episode increased perceived price risk most for cohorts exposed simultaneously to several salient expenditure categories. Respondents born between 1970 and 1980 are more likely to be in peak working and family-forming years, when food, energy, and housing costs are all central to the household budget. Narratives of inflation incidence similarly emphasize that recent price gains in categories such as gasoline and vehicles can weigh heavily on younger and middle-aged households.[21] The stronger increase in education and medical-care uncertainty among older cohorts is also plausible, since these categories can remain financially consequential later in the life cycle through health expenditures and family-related transfers. Overall, the post-COVID patterns reinforce the same point as the cross-sectional gradients: RNRI uncertainty moves systematically with demographic exposure to salient price risks.

[21] A Look at Inflation's Impact by Income and Age in the U.S..

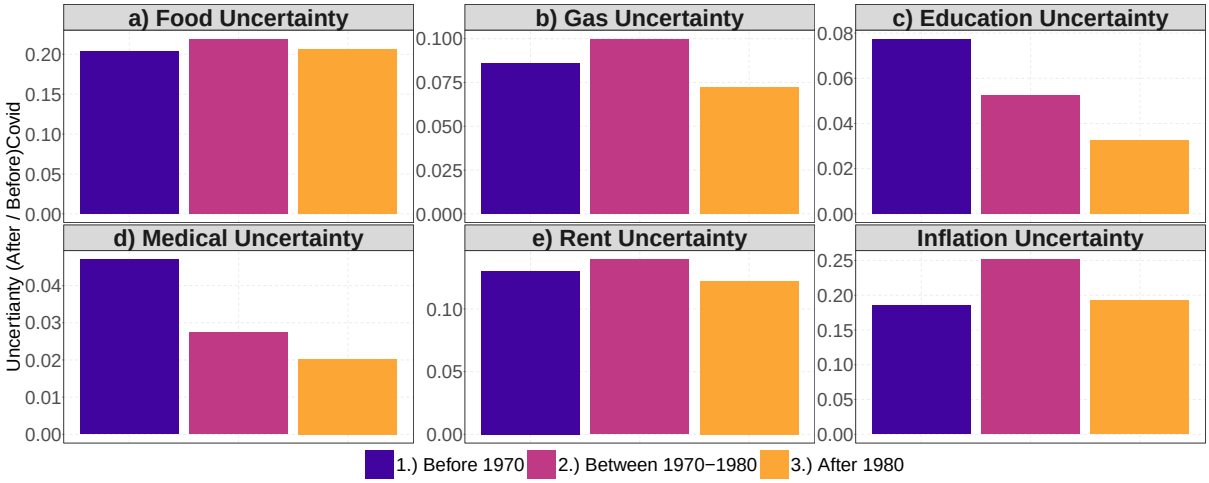


Figure H.10: Age and Commodity Uncertainty Before and After COVID-19

Note: This figure compares sector-specific uncertainty by age group before and after the COVID-19 pandemic (March 2020). Values represent the post-COVID to pre-COVID ratio of uncertainty minus one, (After/Before) − 1, for each age–sector pair. The sample period is June 2013 to June 2025.

I The COVID Window: Event-Study Evidence and Outcome-Date Fixed Effects

March 2020 contains both the largest easing surprise in the sample (the March 3 emergency cut, which carries the entire monthly Jarociński and Karadi (2020) surprise; the March 15 announcement is a zero in the high-frequency series) and the pandemic shock itself. Within-month timing separates the two. Table I.6 reports mean sectoral uncertainty by interview window: uncertainty is essentially flat between the pre-cut window and the two weeks after the March 3 cut, and jumps only in the window beginning March 15, the week of the national emergency declaration and after the March 8 oil-price collapse. The month's uncertainty surge is therefore timed to the lockdown and the oil-price war, not to the monetary easing twelve days earlier; attributing it to the shock would sign a pandemic surprise as a monetary one. Figure I.11 plots the within-month path. The same pandemic and oil-price-war confound continues to operate through the lockdown and reopening months that follow, so rather than drop any observation the baseline projections absorb the acute COVID window (March through June 2020) with outcome-date fixed effects. Concretely, at each horizon h the projection includes indicators for the horizon-h outcomes that fall in these four months, which purge the pandemic and oil-price-war variation from $y_{i,t+h}$ while leaving the rotating SCE panel intact. This treatment is more conservative than simply dropping the window, since it discards no data, and yields the sectoral pattern reported in the main text; the aggregate uncertainty decline is essentially unchanged relative to the drop-the-window variant.

Two further checks harden this treatment. First, rather than absorbing the acute window with outcome-date fixed effects, we drop the March 2020 observation from the shock series entirely, so the largest surprise in the sample contributes no identifying variation at any horizon; the trough responses survive and mostly deepen (Table I.7, Panel A), confirming that the pandemic confound biased the estimates against the measured decline rather than creating it. The three treatments in fact line up monotonically: with no

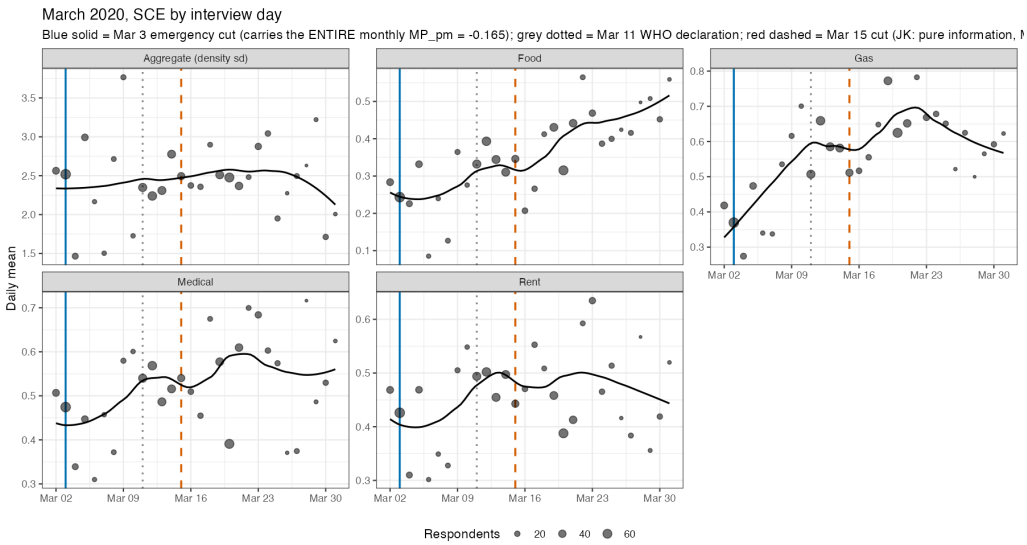


Figure I.11: Within-month path of household uncertainty around March 2020

Note: Mean household uncertainty by interview date within March 2020, from the SCE microdata. The vertical markers indicate the March 11 WHO pandemic declaration and the March 15 emergency declaration and rate cut (a zero in the high-frequency series); the March 3 emergency cut, which carries the entire monthly monetary surprise, precedes the surge shown.

Table I.6: Mean sectoral uncertainty by interview window, February–April 2020

Month	Window	N	Gas	Food	Medical	Rent	Agg. SD
2020-02	d1-2 (pre Mar-3 cut)	63	0.40	0.26	0.59	0.34	2.12
2020-02	d3-14	409	0.38	0.24	0.50	0.38	2.12
2020-02	d15-31	504	0.38	0.24	0.48	0.38	2.02
2020-03	d1-2 (pre Mar-3 cut)	34	0.42	0.28	0.51	0.47	2.56
2020-03	d3-14	400	0.51	0.30	0.49	0.45	2.43
2020-03	d15-31	398	0.63	0.39	0.54	0.46	2.48
2020-04	d1-2 (pre Mar-3 cut)	23	0.69	0.42	0.52	0.42	2.82
2020-04	d3-14	325	0.68	0.47	0.57	0.48	2.64
2020-04	d15-31	391	0.70	0.47	0.58	0.53	2.57

Note: Mean sector-specific uncertainty (χ) and the cross-sectional standard deviation of one-year-ahead point forecasts by within-month interview window in the SCE microdata. The March 3 window pre-dates the emergency cut; the uncertainty surge appears only in the window beginning March 15.

COVID adjustment at all the aggregate uncertainty trough is −0.053 (s.e. 0.023), with the baseline outcome-date fixed effects −0.060 (0.021), and with the March 2020 shock removed −0.067 (0.023), so the finding does not depend on the treatment, and the estimates strengthen as the pandemic observation is removed. Second, placebo projections of pre-shock outcomes on the shock, which should be zero because a monetary surprise cannot move outcomes realized before it, show no pre-trends: across all six uncertainty outcomes and four pre-horizons no coefficient is significantly negative, and the aggregate pre-shock coefficients are essentially zero (Panel B).

Table I.7: Hardening the Uncertainty Local Projections: March 2020 and Placebo Checks

	A. Trough response (per 25bp)			**B. Placebo: pre-shock outcomes**			
Outcome	h	Baseline	Ex-March 2020	$h=-1$	$h=-2$	$h=-3$	$h=-4$
Aggregate	5	-0.060***	-0.067***	0.003	-0.007	-0.003	0.072*
		(0.021)	(0.023)	(0.029)	(0.030)	(0.031)	(0.038)
Gas	6	-0.121	-0.082	-0.076	-0.010	-0.019	-0.066
		(0.077)	(0.088)	(0.063)	(0.072)	(0.066)	(0.066)
Food	5	-0.112**	-0.162***	-0.030	0.019	0.116**	-0.016
		(0.055)	(0.061)	(0.037)	(0.047)	(0.047)	(0.053)
Medical	4	-0.165***	-0.143**	-0.050	0.005	0.026	0.029
		(0.058)	(0.067)	(0.041)	(0.053)	(0.052)	(0.056)
Education	5	-0.044	-0.054	-0.086	0.035	0.012	0.143**
		(0.048)	(0.054)	(0.056)	(0.060)	(0.066)	(0.062)
Rent	5	-0.108**	-0.112**	-0.010	-0.044	0.051	0.026
		(0.045)	(0.048)	(0.043)	(0.050)	(0.051)	(0.039)

Notes: Panel A reports the trough coefficient of each uncertainty local projection (baseline outcome-FE specification of Section 5.1) at the baseline trough horizon h, and the same coefficient when the March 2020 observation is removed from the shock series entirely (rather than absorbed by outcome-date fixed effects). Panel B reports placebo projections of pre-shock outcomes $y_{i,t-h}$ on the shock under the baseline specification; a monetary surprise cannot cause outcomes realized before it, so these coefficients should be zero. Standard errors clustered by survey month in parentheses. Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

J New Keynesian Model with Item-Specific Uncertainty Shocks

This appendix records the full structural model and the derivations the main text defers to it. It states the complete equation system, derives the trend-inflation price-dispersion channel and the exact square-root signal-extraction and welfare blocks, and reports the calibration, moment fit, and determinacy diagnostics. The model is solved as the exact nonlinear Euler, Calvo, and Kalman system under a second-order pruned perturbation, with impulse responses read from the first-order solution; the log-linear uncertainty wedges are an interpretation of these same equations rather than additions to them.

J.1 Relationship to the Literature

Our model sits at the intersection of several strands of work on inflation expectations and monetary policy. The first is the rational-inattention and sticky-information New Keynesian literature, in which information frictions, rather than nominal rigidities alone, slow the response of prices and beliefs to shocks (Sims, 2003; Mankiw and Reis, 2002; Maćkowiak and Wiederholt, 2009, 2015). Like that literature, we treat the precision with which agents perceive sectoral fundamentals as the key friction, and we summarize limited information processing in an estimated dispersion pass-through κ_{disp} that maps the realized cross-item dispersion of price changes into the noise of the price signals households and firms observe, together with a cross-category salience loading through which the most visible dispersion contaminates the processing of other categories. Where we depart is in the source of that dispersion: under positive trend inflation the cross-item dispersion of year-on-year price changes responds to sectoral inflation at first order, so the monetary-policy stance moves the subjective variance around inflation expectations endogenously through its effect on realized price-change dispersion, rather than through a hand-set loading of the policy stance on precision. The second strand is the incomplete-information New Keynesian literature, which uses higher-order belief dynamics and survey-based information rigidities to discipline expectation formation and the slope of the Phillips curve (Coibion and Gorodnichenko, 2012, 2015a; Angeletos and Lian, 2018; Angeletos and Huo, 2021). We share with this work the view that survey expectations are informative about deep parameters of the belief-formation process, but where that literature uses surveys to discipline a scalar information rigidity or a common attention parameter, we use the household-survey impulse responses of uncertainty (Section 5) to estimate the dispersion and salience pass-throughs by minimum distance, with the timing of the second sector's response an over-identifying prediction, complementing the scalar information-rigidity and common-attention calibrations used in that work. A firm-side complement is Flynn et al. (2026), who show that when firms choose supply schedules ex ante under uncertainty about demand and the price level, the optimal schedule is log-linear and its slope steepens endogenously in inflation uncertainty; relative to that mechanism, our fixed-Calvo block holds the pricing slope constant and instead routes the response of inflation uncertainty through the equilibrium dispersion of price changes under trend inflation, so embedding their endogenous-slope channel inside our environment would steepen the pass-through from policy to inflation and hence to dispersion, strengthening rather than offsetting the transmission documented below. A third strand is the multi-sector New Keynesian literature on optimal monetary policy across sectors with heterogeneous price rigidity (Aoki, 2001; Benigno, 2004; Erceg and

Levin, 2006; Mankiw and Reis, 2003). That literature reaches a now-classic prescription: when sectors differ in nominal rigidity, welfare-based optimal policy stabilizes a weighted measure of sticky-price (core) inflation and largely accommodates flexible-price sectors, because the welfare losses from price stickiness are concentrated where prices adjust least. We share its two-sector optimal-policy structure and its heterogeneity in nominal rigidity, with Calvo non-adjustment $\alpha_1 = 0.70$ in food and gasoline against $\alpha_2 = 0.92$ in the remaining basket, but we pose a different question and accordingly reach a different prescription. The margin we isolate is not the stabilization of sticky-price inflation but the trend-inflation dispersion channel: because food and gasoline prices are flexible and highly policy-sensitive, policy moves their inflation, and hence their price-change dispersion and the uncertainty households attach to them, far more than it moves the corresponding objects in services, so our rule leans more on the flexible, demand-sensitive food-and-gasoline sector, in contrast to the lean-on-the-sticky-sector logic of Aoki (2001) and Benigno (2004). The two prescriptions are complementary rather than contradictory: ours isolates the uncertainty-stabilization motive that the sticky-price benchmark holds fixed. A further strand is the older literature on the welfare cost of inflation uncertainty, going back to Friedman (1977) and formalized in Cukierman and Meltzer (1986) and Ball (1992), which establishes that the second moment of inflation, not only its mean, carries first-order welfare consequences. Our two-sector setup operationalizes that insight inside an estimable New Keynesian environment with a micro-founded, time-varying posterior basket-uncertainty object U_t that the central bank can respond to. Two features distinguish our contribution. First, we estimate the transmission parameters directly from the household-survey uncertainty impulse responses rather than from aggregate forecast-error regressions; we view this as a feasible empirical-to-DSGE bridge rather than a set of sharp structural estimates, since it rests on a limited set of matched impulse-response moments (Appendix J.10), with the timing of the services response an over-identifying check. Second, we evaluate a sectorally differentiated uncertainty-response rule $\phi_{U,j}$ that prices the welfare gain of responding more strongly to uncertainty in salient, demand-sensitive sectors (food, gasoline) than in less policy-sensitive ones, a margin that is absent both in single-sector noisy-information NK models and in the sticky-versus-flexible sectoral-policy literature. Together these features connect the rational-inattention micro-foundations of the first strand, the survey-based identification of the second, the sectoral-policy structure of the third, and the welfare logic of the last into a single estimable policy framework.

The "structural noisy-information" equilibrium concept. Throughout the paper "structural noisy-information" refers to the equilibrium concept in which households and firms condition on the posterior of a Bayesian filter, $(m_{j,t}, P_{ij,t})$, rather than on the latent fundamentals $(z_{1,t}, z_{2,t}, z_{d,t})$ directly. At first order the equilibrium policy functions depend only on the posterior means, so Dynare's full-information rational-expectations solution coincides with the strict noisy-information equilibrium. We solve the system at second order with pruning, and at that order the perturbation expansion includes the latent $z_{j,t}$ through quadratic cross-products of the state vector. The order-3 solution should therefore be read as a tractable approximation to the strict noisy-information equilibrium rather than as the equilibrium itself; the welfare and IRF comparisons reported below are robust to this approximation because the dominant policy-sensitive object, the posterior basket-uncertainty gap U_t, is itself a first-order function of the posterior covariances $(P_{11,t}, P_{22,t}, P_{12,t})$.

J.2 Model Setup and Complete Equation System

This subsection records the complete structural system that is solved numerically in Section J.11. Four of its equations are stated in the main text and are cited rather than repeated here: the basket-uncertainty object U_t of Equation (17), the price-change-dispersion recursion (15), the signal-noise state (16), and the two-margin Taylor rule (18). The household and firm optimality conditions, the exact nonlinear consumption Euler equation (J.15) and the recursive Calvo reset-price block (J.21), are stated below within their blocks. The remaining core equations are collected below, grouped by block, so that every structural relation of the model can be found either in the main text or here. The system is solved as the exact nonlinear Euler, Calvo, and Kalman equations under a second-order pruned perturbation; the uncertainty wedges of the following subsection are a log-linear reading of these same equations and are not added to them.

The representative household consumes two sectoral bundles $j \in \{1, 2\}$, with Sector 1 the salient non-durable necessities (food and gasoline) and Sector 2 the rest of the basket. Writing $c_t \equiv \log C_t$ for log aggregate consumption and $C_{j,t}$ for the sectoral bundle, the consumption index and its dual price index are the constant-elasticity aggregators

$$C_t = \left[\omega_1^{1/\eta}\, C_{1,t}^{(\eta-1)/\eta} + (1-\omega_1)^{1/\eta}\, C_{2,t}^{(\eta-1)/\eta}\right]^{\eta/(\eta-1)}, \tag{J.7}$$

$$1 = \omega_1\, e^{(1-\eta)\tilde{p}_{1,t}} + (1-\omega_1)\, e^{(1-\eta)\tilde{p}_{2,t}}, \tag{J.8}$$

where ω_1 is the expenditure share on Sector 1, η the elasticity of substitution across the two bundles, and $\tilde{p}_{j,t} \equiv p_{j,t} - p_t$ the log relative price of Sector j. Equation (J.8) is the exact CES price index written in relative-price form. Within each sector a continuum of varieties aggregates with elasticity θ_j, which fixes the desired gross markup $\mu_j = \theta_j/(\theta_j - 1)$ and governs the Calvo block below. Sectoral and headline inflation are $\pi_{j,t} = \log(P_{j,t}/P_{j,t-1})$ and $\pi_t = \log(P_t/P_{t-1})$, and the relative prices evolve as

$$\tilde{p}_{j,t} = \tilde{p}_{j,t-1} + \pi_{j,t} - \pi_t, \qquad j \in \{1, 2\}. \tag{J.9}$$

There is no separate identity for π_t: headline inflation is pinned jointly by the sectoral rates $(\pi_{1,t}, \pi_{2,t})$ through the relative-price laws (J.9) and the price index (J.8), and reduces to the expenditure-weighted average $\pi_t \approx \omega_1\pi_{1,t} + (1-\omega_1)\pi_{2,t}$ to first order.

The demand-side closure defines the objects that enter the exact Euler equation (J.15). Let $a_{j,t} \equiv \log A_{j,t}$ be sectoral log technology and $\log \xi_t^{\text{eff}} \equiv \log \xi_t + \lambda_d m_{d,t}$ the effective log preference shifter of the main text, in which the posterior mean $m_{d,t}$ of the common demand state enters with loading λ_d. Effective aggregate technology a_t^{eff} is the CES combination of the sectoral technologies,

$$e^{(\eta-1)a_t^{\text{eff}}} = \omega_1\, e^{(\eta-1)a_{1,t}} + (1-\omega_1)\, e^{(\eta-1)a_{2,t}}, \tag{J.10}$$

and the flexible-price (natural) level of log consumption, the output gap, the Wicksellian

natural real rate, and the real-rate gap are

$$c_t^{\text{nat}} = \frac{\log \xi_t^{\text{eff}} - \log \chi_t + (1+\varphi)\, a_t^{\text{eff}}}{\sigma_c + \varphi}, \tag{J.11}$$

$$x_t = c_t - c_t^{\text{nat}}, \tag{J.12}$$

$$r_t^{\text{nat}} = \sigma_c\big(c_{t+1}^{\text{nat}} - c_t^{\text{nat}}\big) - \big(\log \xi_{t+1}^{\text{eff}} - \log \xi_t^{\text{eff}}\big), \tag{J.13}$$

$$s_t = i_t - \mathbb{E}_t \pi_{t+1} - r_t^{\text{nat}}, \tag{J.14}$$

where σ_c is the inverse intertemporal elasticity of substitution, φ the inverse Frisch elasticity, $\log \chi_t$ the labor-disutility shifter, i_t the log gross nominal rate, and x_t and s_t the output gap and real-rate gap that enter the policy rule (18). The natural rate (J.13) is the real rate that would obtain under flexible prices and drives s_t to zero absent frictions.

On the demand side, the household's consumption–savings first-order condition is the exact nonlinear Euler equation

$$1 = \exp(\tilde{i}_t)\, \mathbb{E}_t\left[\exp\left(\log \frac{\xi_{t+1}^{\text{eff}}}{\xi_t^{\text{eff}}} - \sigma_c \log \frac{C_{t+1}}{C_t} - \pi_{t+1}\right)\right], \tag{J.15}$$

where $\tilde{i}_t \equiv \log R_t + \log \beta$ is the steady-state-normalized log gross nominal interest rate, so the household discount factor β is absorbed into $\tilde{i}_t$. At a log-linear approximation the elasticity of expected consumption growth to the real rate is $1/\sigma_c$, so σ_c governs the strength of the demand channel; the effective preference shifter ξ_t^{eff} defined above brings the posterior mean of the common demand-pressure state into the same condition.[22]

On the production side, CES demand splits consumption across sectors and hours are demanded to meet it net of the Calvo misallocation wedge. Writing $c_{j,t} \equiv \log C_{j,t}$ and $n_t \equiv \log N_t$ for log hours,

$$c_{j,t} = c_t - \eta\, \tilde{p}_{j,t}, \qquad j \in \{1,2\}, \tag{J.16}$$

$$n_{j,t} = c_{j,t} + \log\big(1 + \Delta_{j,t}\big) - a_{j,t}, \qquad j \in \{1,2\}, \tag{J.17}$$

$$e^{n_t} = \omega_1\, e^{n_{1,t}} + (1-\omega_1)\, e^{n_{2,t}}, \tag{J.18}$$

where $\Delta_{j,t}$ is the Calvo measure of within-sector price dispersion whose law of motion is Equation (J.24) below; the term $\log(1+\Delta_{j,t})$ is the extra labor absorbed by relative-price dispersion, the channel through which sectoral misallocation enters hours and hence welfare. The real wage follows from the household intratemporal condition, and sectoral log real marginal cost from cost minimization,

$$w_t = \log \chi_t - \log \xi_t^{\text{eff}} + \sigma_c c_t + \varphi\, n_t, \tag{J.19}$$

$$mc_{j,t} = w_t - a_{j,t} - \tilde{p}_{j,t}, \qquad j \in \{1,2\}. \tag{J.20}$$

The effective marginal cost $mc_{j,t}^{\text{eff}} = \exp(mc_{j,t} + \lambda_{u,j} m_{j,t})/\mu_j$ that enters the reset-price numerator, in which the firm's posterior mean $m_{j,t}$ of sectoral markup pressure loads with weight $\lambda_{u,j}$, enters the numerator recursion of Equation (J.21) below. Goods-market

[22] The implementation allows the loading to combine the posterior mean with the underlying fundamental $z_{d,t}$ as $\lambda_d\big(\theta_I\, m_{d,t} + (1-\theta_I)\, z_{d,t}\big)$, with information share $\theta_I \in [0,1]$. The baseline sets $\theta_I = 1$, so only the posterior mean enters.

clearing is implicit: sectoral output equals the sectoral demand (J.16), and total hours are the aggregate (J.18).

The Calvo block operates under positive trend inflation $\bar{\pi} = \log(1.02)/4$, so the pricing recursions carry the trend gross inflation rate and the price level is dispersed even in steady state. The recursive numerator $G_{j,t}$ and denominator $F_{j,t}$ of the optimal reset price follow

$$\begin{aligned} G_{j,t} &= C_{j,t}\, mc^{\text{eff}}_{j,t} + \alpha_j \beta\, \mathbb{E}_t \left[\exp\left(\log \frac{\xi^{\text{eff}}_{t+1}}{\xi^{\text{eff}}_t} - \sigma_c \log \frac{C_{t+1}}{C_t} + \theta_j \pi_{j,t+1} \right) G_{j,t+1} \right], \\ F_{j,t} &= C_{j,t} + \alpha_j \beta\, \mathbb{E}_t \left[\exp\left(\log \frac{\xi^{\text{eff}}_{t+1}}{\xi^{\text{eff}}_t} - \sigma_c \log \frac{C_{t+1}}{C_t} + (\theta_j - 1) \pi_{j,t+1} \right) F_{j,t+1} \right], \end{aligned} \tag{J.21}$$

Writing $p^*_{j,t} \equiv \log(P^*_{j,t}/P_{j,t})$ for the log reset price relative to the sectoral price level, that block delivers the reset-price condition and the Calvo inflation identity

$$e^{p^*_{j,t}} = \mu_j \frac{G_{j,t}}{F_{j,t}}, \qquad j \in \{1,2\}, \tag{J.22}$$

$$1 = \alpha_j\, e^{(\theta_j - 1)\pi_{j,t}} + (1 - \alpha_j)\, e^{(1-\theta_j) p^*_{j,t}}, \qquad j \in \{1,2\}, \tag{J.23}$$

where α_j is the Calvo non-adjustment probability. Equation (J.23) is the standard price-index restriction that links sectoral inflation to the reset price, a fraction $1 - \alpha_j$ of firms resetting to $p^*_{j,t}$ each period while the rest carry their price forward. The associated law of motion for within-sector relative-price dispersion is the Yun recursion

$$\Delta_{j,t} = \alpha_j\, e^{\theta_j \pi_{j,t}} \left(\Delta_{j,t-1} + 1 \right) + (1 - \alpha_j)\, e^{-\theta_j p^*_{j,t}} - 1, \qquad j \in \{1,2\}, \tag{J.24}$$

which closes the labor block (J.17). This price-level dispersion $\Delta_{j,t}$, the source of misallocation in hours, is distinct from the cross-item variance of price changes $\left(V_{j,t}, D^{\text{YoY}}_{j,t} \right)$ of Equation (15), which is the signal noise households sample when forming beliefs.

The belief block that turns dispersion into posterior uncertainty is the mechanism of the paper and is documented separately: the mapping from the equilibrium dispersion (15) through the signal-noise state (16) to the pass-through κ is set out in Appendix J.3, and the exact square-root (Cholesky) Kalman filter that produces the posterior means $(m_{1,t}, m_{2,t}, m_{d,t})$, the posterior covariance, the sectoral inflation-forecast variances $\sigma^2_{\pi,j,t}$, their correlation ρ_t, and the basket object U_t of Equation (17) is in Appendix J.5.

The exogenous block drives the system through the latent states that the belief block filters. Sectoral markup pressure and the common demand state follow

$$z_{1,t} = \rho_z\, z_{1,t-1} + \sigma_z\, \varepsilon^{z_1}_t, \tag{J.25}$$

$$z_{2,t} = \rho_z\, z_{2,t-1} + \sigma_z \left(\text{corr}_z\, \varepsilon^{z_1}_t + \sqrt{1 - \text{corr}_z^2}\, \varepsilon^{z_2}_t \right), \tag{J.26}$$

$$z_{d,t} = \rho_{zd}\, z_{d,t-1} + \sigma_{zd}\, \varepsilon^{z_d}_t, \tag{J.27}$$

with $\rho_z = 0.55$, $\sigma_z = 0.118$, $\rho_{zd} = 0.60$, and $\sigma_{zd} = 0.040$. The sector-2 markup innovation in (J.26) loads on the sector-1 shock $\varepsilon^{z_1}_t$ with weight $\text{corr}_z = 0.055$: this correlation is the only structural coupling between the two sectors' fundamentals, and it is what generates the model's realized cross-sector inflation comovement, all other cross-sector

linkage running through the common demand state $z_{d,t}$ and the shared belief filter. The preference, labor, and sectoral technology shifters follow stationary autoregressions,

$$\log \xi_t = (1-\rho_\xi)\bar{\xi} + \rho_\xi \log \xi_{t-1} + \sigma_\xi \varepsilon_t^{\xi}, \tag{J.28}$$

$$\log \chi_t = (1-\rho_\chi)\bar{\chi} + \rho_\chi \log \chi_{t-1} + \sigma_\chi \varepsilon_t^{\chi}, \tag{J.29}$$

$$a_{j,t} = (1-\rho_A)\bar{a}_j + \rho_A\, a_{j,t-1} + \sigma_A\, \varepsilon_t^{A_j}, \qquad j \in \{1,2\}, \tag{J.30}$$

whose innovation standard deviations are set to zero in the baseline, $\sigma_\xi = \sigma_\chi = \sigma_A = 0$, so that the baseline is driven by the markup, demand, monetary, observation-noise, and dispersion shocks alone; the processes (J.28)–(J.30) are retained for the robustness analysis. The observation-noise and independent-dispersion shocks that also enter the innovations are described with the belief block in Appendix J.5.

Two objects are computed for diagnostics and feed no other equation. The subjective posterior-mean forecast of cumulative four-quarter Sector-j inflation,

$$E^{\pi}_{j,t} = \lambda^{\pi_j}_{m_1} m_{1,t} + \lambda^{\pi_j}_{m_2} m_{2,t} + \lambda^{\pi_j}_{m_d} m_{d,t}, \tag{J.31}$$

where the loadings $\left(\lambda^{\pi_j}_{m_1}, \lambda^{\pi_j}_{m_2}, \lambda^{\pi_j}_{m_d}\right)$ are the model-implied cumulative-inflation weights on the posterior means, records the first moment of beliefs about the exogenous sectoral fundamentals; because those latent states are exogenous to policy, this diagnostic has no first-order response to a policy shock. It is distinct from households' forward-looking inflation expectations, which do respond to the policy shock through the standard channel and match the survey expectations local projections (Figure J.14). The augmented rule responds to the posterior inflation variance, a second moment that moves alongside, rather than in place of, that conventional first-moment response. The alternative price-change-dispersion measure

$$Dpc_{j,t} = \frac{\alpha_j}{1-\alpha_j}\pi^2_{j,t} + (1-\alpha_j)\frac{2}{\theta_j^2}\,\Delta_{j,t-1} \tag{J.32}$$

combines a transitory inflation term with the carried price-level dispersion of Equation (J.24); it is reported for comparison only, the live signal-noise channel running through $\left(V_{j,t}, D^{\text{YoY}}_{j,t}\right)$ in Equation (15).

J.3 The Trend-Inflation Dispersion Channel

The object that governs how monetary policy moves sectoral inflation uncertainty is not a hand-set coefficient on the policy stance but a solved equilibrium chain running from inflation, through the cross-item dispersion of price changes, into the observation noise of the households' and firms' signal-extraction problem. This subsection derives that chain from the Calvo cross-section, shows that positive trend inflation $\bar{\pi} \equiv \log(1.02)/4 = 0.00495$ is precisely what makes it operate at first order, and states the resulting policy semi-elasticity of uncertainty as a derived rather than an imposed quantity.

Fix a sector j and let $p^*_{j,t}$ denote the common (log) reset price chosen by every sector-j firm that re-optimizes at t. Under Calvo pricing a firm re-optimizes with probability $1-\alpha_j$ and otherwise carries its previous price, where $\alpha_j \in (0,1)$ is the non-adjustment probability ($\alpha_1 = 0.70$ for the flexible food and gasoline sector, $\alpha_2 = 0.92$ for services). Write the

gap between item i's posted log price and the current reset price as $d_{i,t} \equiv p_{i,t} - p^*_{j,t}$. A firm drawn to re-optimize posts $d_{i,t} = 0$; a firm that keeps its price carries

$$d_{i,t} = d_{i,t-1} - \Delta p^*_{j,t}, \qquad \Delta p^*_{j,t} = \pi_{j,t} + O(\|\cdot\|^2), \tag{J.33}$$

because along the trend-inflation path the reset price and the sectoral price level both drift at the sectoral inflation rate $\pi_{j,t} = \log(P_{j,t}/P_{j,t-1})$ to first order. Since the Calvo draw is independent of an item's vintage, the cross-sectional mean gap $\mu_{j,t} \equiv \mathbb{E}_i(d_{i,t})$ and variance $V_{j,t} \equiv \mathrm{Var}_i(d_{i,t})$ obey

$$\mu_{j,t} = \alpha_j\Big(\mu_{j,t-1} - \pi_{j,t}\Big), \tag{J.34}$$

$$V_{j,t} = \alpha_j V_{j,t-1} + \alpha_j(1-\alpha_j)\Big(\mu_{j,t-1} - \pi_{j,t}\Big)^2, \tag{J.35}$$

where the second line follows from expanding $V_{j,t} = \mathbb{E}_i[d^2_{i,t}] - \mu^2_{j,t}$ using the two-point mixture over resetters and non-resetters and collecting the mean cross-term. Along the balanced path the mean gap is $\bar{\mu}_j = -\frac{\alpha_j}{1-\alpha_j}\bar{\pi}$, so that in steady state $\mu_{j,t-1} - \pi_{j,t} = -\bar{\pi}/(1-\alpha_j)$. The live model carries the exact system (J.35), with the mean reset-price gap $\mu_{j,t}$ as a state; it is the sectoral form of the main-text recursion (15). Substituting the steady-state mean gap into the innovation term would collapse it to $\frac{\alpha_j}{1-\alpha_j}\pi^2_{j,t}$; that substitution is exact only along the deterministic trend and front-loads the dispersion response around it, so the model does not use it.

The empirical counterpart to household price signals is the cross-item dispersion of year-on-year price changes rather than of price levels. Item i's year-on-year log price change is $p_{i,t} - p_{i,t-4}$, whose cross-item variance is $\mathrm{Var}_i(p_{i,t}) + \mathrm{Var}_i(p_{i,t-4}) - 2\,\mathrm{Cov}_i(p_{i,t}, p_{i,t-4})$. The level dispersion is $V_{j,t}$ net of the common trend, and to a close approximation the only items with a positive cross-period covariance are the fraction α^4_j that have not re-optimized in four quarters, for which $p_{i,t} = p_{i,t-4}$; this gives $\mathrm{Cov}_i(p_{i,t}, p_{i,t-4}) \approx \alpha^4_j V_{j,t-4}$ and hence

$$D^{YoY}_{j,t} = V_{j,t} + \Big(1 - 2\alpha^4_j\Big)V_{j,t-4} + \sigma_{Dj}\,\varepsilon^D_{j,t}, \tag{J.36}$$

the year-on-year price-change dispersion of the main text. The last term is an independent dispersion shock, discussed below, which is switched off for the derivation that follows.

The central result is that with $\bar{\pi} > 0$ dispersion loads on inflation at first order, the trend-inflation mechanism of Ascari (2004) and Ascari and Sbordone (2014) put to work as an information channel. The system (J.35) has steady state $\bar{V}_j = \frac{\alpha_j}{(1-\alpha_j)^2}\bar{\pi}^2$; linearizing around trend, with $\hat{\mu}_{j,t}$ and $\hat{V}_{j,t}$ the deviations of the mean gap and the variance and $\hat{\pi}_{j,t} \equiv \pi_{j,t} - \bar{\pi}$,

$$\hat{\mu}_{j,t} = \alpha_j\Big(\hat{\mu}_{j,t-1} - \hat{\pi}_{j,t}\Big), \qquad \hat{V}_{j,t} = \alpha_j\hat{V}_{j,t-1} + 2\alpha_j\bar{\pi}\Big(\hat{\pi}_{j,t} - \hat{\mu}_{j,t-1}\Big), \qquad \left.\frac{\partial D^{YoY}_{j,t}}{\partial \pi_{j,t}}\right|_{ss} = 2\alpha_j\bar{\pi}, \tag{J.37}$$

since the predetermined lags contribute nothing to the contemporaneous derivative. The impact loading of dispersion on inflation is therefore $2\alpha_j\bar{\pi}$, equal to 0.0069 for food and gasoline and 0.0091 for services, and the mean-gap state propagates the impulse forward: a one-off inflation innovation keeps feeding dispersion through $\hat{\mu}_{j,t}$, with cumulative direct loading $2\alpha_j\bar{\pi}/(1-\alpha_j)$, so the dispersion response builds over several quarters rather than arriving at once. At $\bar{\pi} = 0$ both pieces are identically zero: the surviving term in $\pi^2_{j,t}$

is the second-order piece $\hat{\pi}^2_{j,t}$, so dispersion would respond to inflation only at second order and the channel would be inert in a first-order reading of the model. Positive trend inflation is what activates the mechanism at first order, and is the structural primitive from which the entire policy-to-uncertainty transmission is derived.

The impact loading is nearly common across sectors, so the cross-sector asymmetry is carried by the sectoral inflation response itself. The food/gas Phillips curve is steep, with slope $\kappa_1 = 0.132$, whereas the services curve is nearly flat, with $\kappa_2 = 0.008$, a ratio of about sixteen: a contractionary shock moves food and gasoline inflation far more than services inflation, so the net policy-to-dispersion-to-uncertainty effect is far stronger in food and gasoline. This is the structural source of the cross-sector impulse-response asymmetry: the model's services-to-food/gas uncertainty trough ratio is about 0.42, close to the 0.40 ratio in the data.

Households and firms do not observe the latent sectoral markup state $z_{j,t}$ directly; they read it from noisy price data, and the informativeness of those data is governed by how dispersed item price changes are. Write the noisy price signal as $\hat{s}_{j,t} = z_{j,t} + \sqrt{R_{j,t}}\,\varepsilon^{\text{sig}}_{j,t}$, where $\varepsilon^{\text{sig}}_{j,t}$ is standard observation noise and the noise variance is

$$\begin{aligned} R_{j,t} &= \bar{R}_j \exp(q_{j,t})\,, \\ q_{j,t} &= \rho_q q_{j,t-1} + \kappa_{\text{disp}}\, \widetilde{D}^{\text{YoY}}_{j,t} \\ &\quad + \mathbf{1}\{j=2\}\left(\kappa_{\text{serv}}\, \widetilde{D}^{\text{YoY}}_{1,t} + \sigma_{h,2}\varepsilon^{h}_{2,t}\right), \\ \widetilde{D}^{\text{YoY}}_{j,t} &\equiv \frac{D^{\text{YoY}}_{j,t} - \bar{D}^{\text{YoY}}_{j}}{\bar{D}^{\text{YoY}}_{j}}. \end{aligned} \tag{J.38}$$

the signal-noise state of the main text (16), with steady-state noise variances $\bar{R}_1 = \bar{R}_2 = 0.02$ and $\bar{R}_d = 0.002$. Services signal noise carries the salience loading κ_{serv} on the food-and-gas dispersion input and the idiosyncratic household-uncertainty innovation $\sigma_{h,2}\varepsilon^h_{2,t}$; food-and-gas signal noise carries neither. The memory $\rho_q = 0.55$ is estimated jointly with the pass-throughs against the persistence of the measured uncertainty index (0.89 against 0.88); so $q_{j,t}$ is a persistent function of realized dispersion. The estimated dispersion pass-through is $\kappa_{\text{disp}} = 2.08$, distinct from the Phillips slopes κ_1, κ_2, and the estimated salience loading is $\kappa_{\text{serv}} = 0.83$. Its interpretation is a limited-information one: an agent forming a belief about sector-j price pressure from n_{eff} effectively independent item-price observations has an observation-noise variance that falls with n_{eff} and rises with the cross-item dispersion of those observations, so that κ_{disp} combines the inverse of the effective number of independently processed item-price signals with a measurement mapping from the model's dispersion units to the scale of the survey uncertainty index. It is therefore an estimated pass-through rather than a structural elasticity bounded below one, and its magnitude reflects that units conversion. A single κ_{disp} is shared across sectors, so all sectoral heterogeneity in the uncertainty response runs through the loadings $2\alpha_j\bar{\pi}/(1-\alpha_j)$ and the Phillips slopes, not through a sector-specific pass-through. The independent dispersion shock $\varepsilon^D_{j,t}$ in (J.36), calibrated at $\sigma_{D1} = 1.45 \times 10^{-4}$ and $\sigma_{D2} = 9.99 \times 10^{-4}$ (a spread of roughly half of $\bar{D}^{YoY}_j$), shifts the cross-item spread without moving the sectoral mean, so uncertainty can move orthogonally to the price level. Together with the accumulation of dispersion through the mean-gap state, which decouples the posterior variance from current inflation, this is why the model-implied

correlation between food/gas uncertainty and food/gas inflation is essentially zero, about -0.11 under a rule that does not respond to uncertainty.

Given $R_{j,t}$, the posterior is produced by the exact square-root (Cholesky) Kalman filter of Appendix J.5, which delivers the posterior belief $m_{j,t}$ and the posterior covariance P_t of $(z_{1,t}, z_{2,t}, z_{d,t})$. The posterior variance of cumulative four-quarter sectoral inflation is the quadratic form $\sigma^2_{\pi,j,t} = a_j' P_t a_j$, where $a_j = (\lambda_{\pi_j,m_1}, \lambda_{\pi_j,m_2}, \lambda_{\pi_j,m_d})'$ collects the cumulative-inflation loadings on the posterior means, and reported uncertainty is its log, $h_{j,t} = \log \sigma^2_{\pi,j,t}$, entering the Taylor rule (18) as the gap $h^{\text{pol}}_{j,t} = h_{j,t} - \bar{h}_j$.

The policy semi-elasticity of sectoral uncertainty is not a free parameter but a solved object of the transmission chain

$$\text{policy} \longrightarrow \pi_{j,t} \longrightarrow D^{YoY}_{j,t} \longrightarrow q_{j,t} \longrightarrow R_{j,t} \longrightarrow P_t \longrightarrow \sigma^2_{\pi,j,t} \longrightarrow h_{j,t}, \quad \text{(J.39)}$$

in which the only new structural parameter is κ_{disp} and every other link is the derivative of an equilibrium object. Locally,

$$\left.\frac{\partial h_{j,t}}{\partial \pi_{j,t}}\right|_{ss} = \underbrace{\frac{\partial h_{j,t}}{\partial q_{j,t}}}_{\text{Kalman, } >0} \quad \underbrace{\frac{\partial q_{j,t}}{\partial D^{YoY}_{j,t}}}_{=\kappa_{\text{disp}}/\bar{D}^{YoY}_j} \quad \underbrace{\frac{\partial D^{YoY}_{j,t}}{\partial \pi_{j,t}}}_{=2\alpha_j\bar{\pi}/(1-\alpha_j)}, \quad \text{(J.40)}$$

where $\partial h_{j,t}/\partial q_{j,t} > 0$ because a higher $q_{j,t}$ raises the observation-noise variance $R_{j,t}$, lowers the Kalman gain, and therefore raises the posterior variance $\sigma^2_{\pi,j,t} = a_j' P_t a_j$; the remaining link $\partial \pi_{j,t}/\partial(\text{policy})$ is the sectoral inflation response to the rule, which the model solves jointly with the belief block. A contractionary innovation lowers $\pi_{j,t}$, hence $D^{YoY}_{j,t}$, $q_{j,t}$ and $R_{j,t}$, sharpening the price signal and lowering $h_{j,t}$, so uncertainty falls after a tightening, reproducing the sign of the empirical impulse responses. Because $\partial h_{j,t}/\partial q_{j,t}$ runs through the nonlinear Kalman recursion, the composite semi-elasticity in (J.40) has no closed form and is reported as a solved-model object. The estimation disciplines κ_{disp} so that the food/gas uncertainty response matches its empirical local-projection counterpart (matching the covariance-aware cumulative three-quarter response to within ten percent, within the 90% band at every quarter), and the salience loading κ_{serv} so that the services trough size matches; the timing and shape of the services response, which import the food-and-gas dispersion dynamics with no parameter of their own, are then predictions of the same mechanism.

J.4 The Measured-Dispersion Block: Gas–Food Relative Price, Volatility Response, and Deferred Adjustment

The BLS dispersion measure of Section 5.2 is the cross-item variance of year-on-year inflation within the non-core basket, and Figure J.13 decomposes it into a within-food component and a between component, the count-share-weighted squared gap between the gas and food group means. This appendix states the model block that carries both components, its discipline, and the verification that it leaves every previously matched object unchanged.

The block adds five equations, all in the measurement layer. The gas–food relative price is driven by a slow energy-demand state and exogenous oil innovations whose volatility

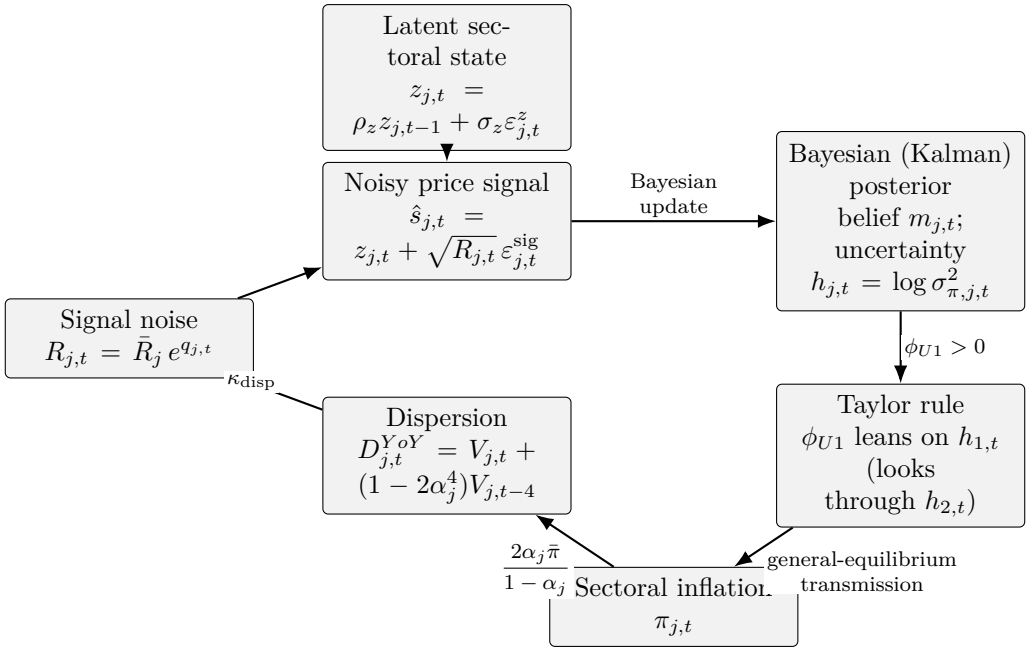


Figure J.12: Information process and the policy-dispersion feedback loop. Households and firms read the latent sectoral state $z_{j,t}$ off noisy price data; the observation-noise variance $R_{j,t} = \bar{R}_j e^{q_{j,t}}$ is set by the cross-item dispersion of year-on-year price changes $D^{YoY}_{j,t}$, with dispersion pass-through $\kappa_{\text{disp}} = 2.08$ and salience loading $\kappa_{\text{serv}} = 0.83$. The Bayesian (Kalman) posterior delivers the belief $m_{j,t}$ and the log posterior inflation-forecast variance $h_{j,t} = \log \sigma^2_{\pi,j,t}$, on which the rule leans through ϕ_{U1} for food and gasoline while looking through services ($\phi_{U2} = 0$). Policy feeds back into sectoral inflation $\pi_{j,t}$ and hence into dispersion with first-order loading $2\alpha_j\bar{\pi}/(1-\alpha_j)$. The asymmetry is structural: the loading is about five times larger in services (0.114) than in food and gasoline (0.023), but a contractionary shock moves the flexible food/gas sector's inflation far more (Phillips slopes $\kappa_1 = 0.132$ versus $\kappa_2 = 0.008$), so the net uncertainty response is larger for food and gasoline. At $\bar{\pi} = 0$ the loading vanishes and the loop is inactive at first order.

falls with the policy stance,

$$d^g_t = \rho_d\, d^g_{t-1} + (1-\rho_d)\,\zeta_g\,(x_t - \bar{x}) + \sigma_g\Big(1 - \psi_v\,(i_t - \bar{\imath})\Big)\varepsilon^g_t, \qquad g_t = d^g_t - d^g_{t-4},$$

where g_t is the four-quarter relative-price change. The slow adjustment is the vehicle-stock margin of gasoline demand, whose short-run price elasticity is an order of magnitude below its long-run value (Hughes et al., 2008); the volatility response captures calmer energy-price innovations under tighter policy. Deferred price adjustment is carried by a pent-up stock,

$$b_t = \rho_b\, b_{t-1} + (i_t - \bar{\imath}),$$

the first-order approximation to the selection dynamics of state-dependent pricing: fewer items reach their adjustment trigger while desired increases are compressed, and the deferred adjustments release in a bunched wave afterwards (Caballero and Engel, 2007; Golosov and Lucas, 2007). Measured non-core dispersion is then

$$D^{\text{meas}}_{1,t} = D^{YoY}_{1,t} + b_{g2}\, g^2_t + \nu_c\, b_{t-1} - \nu_d\,(i_t - \bar{\imath}),$$

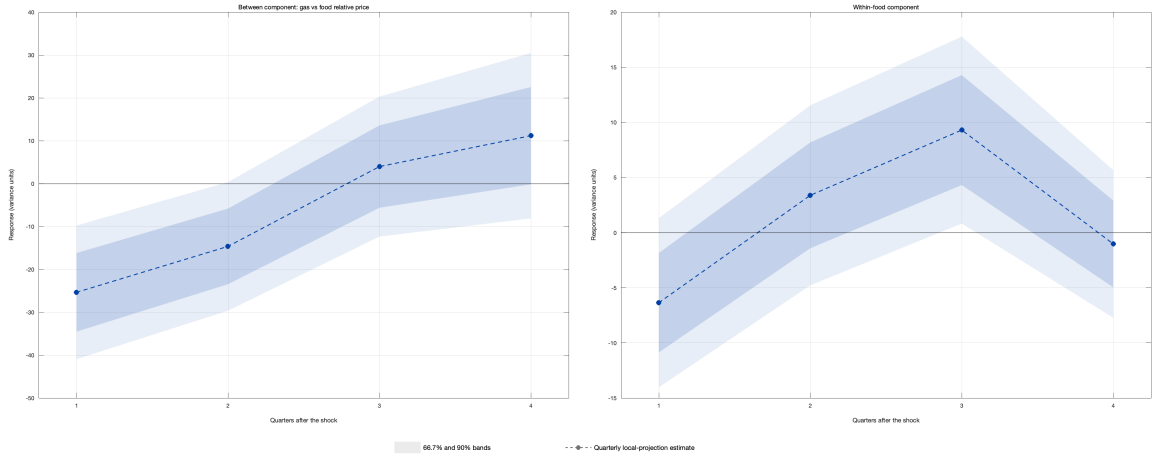


Figure J.13: Decomposing the Non-Core Dispersion Response: Gas versus Food

Note: Variance decomposition of the non-core dispersion response in Figure 10: total cross-item variance of year-on-year inflation within the food-and-gas basket equals the pooled within-group variance plus the between component, the count-share-weighted squared gap between the gas and food group means. Each panel reports quarterly-aggregated local projections of one component on the same specification and sample as Figure 10 (the BLS projections exclude March 2020; as aggregate item-index regressions they carry no panel structure, so exclusion is innocuous there, in contrast to the household-panel treatment of Appendix I); shaded areas are 66.7% and 90% bands. Under a joint moving-block bootstrap across the item panel, the between component's impact compression is significant at the 90% level ($t = -1.8$); the within-food component's third-quarter rebound is the strongest positive estimate in the decomposition but imprecise ($t \approx 1.1$).

and the signal-noise state reads its unanticipated part,

$$q_{1,t} = \rho_q \, q_{1,t-1} + \kappa_{\text{disp}} \frac{D_{1,t}^{YoY} - \bar{D}_1 + b_{g2}\left(g_t^2 - v_t\right)}{\bar{D}_1}, \qquad v_t = \rho_v \, v_{t-1} + (1 - \rho_v) \, g_t^2,$$

with $\rho_v = 0.95$ fixed a priori, a multi-year normalization horizon for the between-term reference level (the inertness verification at the end of this subsection covers the between block). Two exclusions are deliberate. The deferral–catch-up terms do not enter $q_{1,t}$: the wave is predictable given the policy path, so a Bayesian filter discounts it and it does not add signal noise. And the between term enters relative to its slow-moving average v_t, which has zero steady state and removes the term's ergodic mean, matching the standardized survey uncertainty measure while leaving the deterministic steady state and the stochastic means of q_1 and h_1 unshifted.

The block's parameters are estimated jointly with κ_{disp} and ρ_q (Appendix J.13), subject to two measured constants: the gap's unconditional volatility (18pp) is imposed as an identity, and b_{g2} is the measured between-to-within steady-state variance ratio (0.757). At the estimate the gap loads entirely on exogenous oil innovations, $\rho_d = 0.20$ and $\sigma_g = 0.125$, with the activity share estimated at zero; profiling that share confirms zero is preferred even when the gap moments are scored on shape alone, so the policy-conditional between compression runs through the volatility margin: ψ_v is an interior estimate, well inside the admissibility region $1 - \psi_v(i_t - \bar{\imath}) > 0$. The deferral–catch-up pair is $\rho_b = 0.58$ with the catch-up share at its intensity cap, placing the sign change of the dispersion response after the second quarter. Component magnitudes are matched proportionally, not jointly: the between channel's ergodic response cannot reach the episode-leveraged magnitudes of the 2022–2024 tightening, and the units mapping absorbs the difference

exactly as κ_{disp} absorbs the survey scale.

Because the between and deferral components are respectively quadratic and anticipated, they contribute nothing at first order: the sectoral uncertainty impulse responses, the κ_{disp} calibration, the persistence and orthogonality diagnostics, and the entire welfare exercise are numerically identical with and without the block, which we verify directly (the welfare grid agrees point by point to four decimals). The dispersion exhibit in Figure 11 is therefore the model's mean response at second order; the quadratic between contribution is computed by window-averaging over the model's ergodic gap distribution, since simulated generalized impulse responses of squared states carry sampling noise of the same order as the object at feasible replication counts.

J.5 Exact Square-Root Signal Extraction

Households and firms do not observe the latent drivers of inflation directly. The sectoral markup pressures $z_{1,t}$ (food and gasoline) and $z_{2,t}$ (services) and a common demand factor $z_{d,t}$ follow the first-order autoregressions

$$z_{1,t} = \rho_z \, z_{1,t-1} + \sigma_z \, \varepsilon_t^{z1}, \tag{J.41}$$

$$z_{2,t} = \rho_z \, z_{2,t-1} + \sigma_z \left(\text{corr}_z \, \varepsilon_t^{z1} + \sqrt{1 - \text{corr}_z^2} \, \varepsilon_t^{z2} \right), \tag{J.42}$$

$$z_{d,t} = \rho_{zd} \, z_{d,t-1} + \sigma_{zd} \, \varepsilon_t^{zd}, \tag{J.43}$$

where ρ_z and σ_z are the common persistence and innovation standard deviation of the sectoral markup states, ρ_{zd} and σ_{zd} their demand counterparts, and $\{\varepsilon_t^{z1}, \varepsilon_t^{z2}, \varepsilon_t^{zd}\}$ are mutually independent standard-normal innovations. The scalar $\text{corr}_z \in [0, 1)$ is the contemporaneous correlation of the two sectoral markup innovations and is the only channel through which the sectors are coupled at the level of fundamentals. Agents observe noisy price signals of each state,

$$y_{j,t} = z_{j,t} + \sqrt{R_{j,t}} \, \varepsilon_{j,t}^{sig}, \qquad j \in \{1, 2, d\}, \tag{J.44}$$

with $\varepsilon_{j,t}^{sig}$ independent standard normal and $R_{j,t}$ the signal-noise variance, which is time-varying and is supplied by the trend-inflation dispersion channel discussed at the end of this subsection. From these signals agents form posterior means $m_{1,t}, m_{2,t}, m_{d,t}$ of the latent states and a posterior covariance P_t of $(z_{1,t}, z_{2,t}, z_{d,t})$. Because the markup and demand innovations are orthogonal and the demand signal is observed separately, P_t is block diagonal, with a 2×2 sectoral block $(P_{11,t}, P_{12,t}, P_{22,t})$ and a scalar demand entry $P_{dd,t}$.

The filter is carried in square-root (Cholesky) form. Rather than propagate the covariance entries, the model stores the Cholesky factors of P_t as the states $cL1_t, cL2_t, cL3_t$ (sectoral block) and cLd_t (demand), with the diagonal factors held in logs so that they cannot cross zero. The one-step-ahead prior mean of each state is its autoregressive continuation of last period's posterior mean,

$$\hat{a}_{j,t} = \rho_z \, m_{j,t-1} \;\; (j \in \{1, 2\}), \qquad \hat{a}_{d,t} = \rho_{zd} \, m_{d,t-1}. \tag{J.45}$$

The lagged posterior covariance of the sectoral block is reconstructed from the lagged

Cholesky states, $P^{+}_{t-1} = L_{t-1}L_{t-1}^{\top}$ with

$$L_{t-1} = \begin{pmatrix} e^{cL1_{t-1}} & 0 \\ cL2_{t-1} & e^{cL3_{t-1}} \end{pmatrix}, \qquad \begin{aligned} P^{+}_{11,t-1} &= e^{2cL1_{t-1}}, \\ P^{+}_{12,t-1} &= e^{cL1_{t-1}} cL2_{t-1}, \\ P^{+}_{22,t-1} &= cL2^{2}_{t-1} + e^{2cL3_{t-1}}. \end{aligned} \tag{J.46}$$

Propagating this block through the common autoregressive root ρ_z and adding the markup innovation covariance $\sigma_z^2 \left(\begin{smallmatrix} 1 & \mathrm{corr}_z \\ \mathrm{corr}_z & 1 \end{smallmatrix}\right)$ gives the prior (predicted) covariance

$$P^{-}_{11,t} = \rho_z^2 P^{+}_{11,t-1} + \sigma_z^2, \quad P^{-}_{12,t} = \rho_z^2 P^{+}_{12,t-1} + \mathrm{corr}_z\, \sigma_z^2, \quad P^{-}_{22,t} = \rho_z^2 P^{+}_{22,t-1} + \sigma_z^2, \tag{J.47}$$

whose Cholesky factor $(\ell^{-}_1, \ell^{-}_2, \ell^{-}_3)$ is

$$\ell^{-}_1 = \sqrt{P^{-}_{11,t}}, \qquad \ell^{-}_2 = \frac{P^{-}_{12,t}}{\ell^{-}_1}, \qquad \ell^{-}_3 = \sqrt{P^{-}_{22,t} - (\ell^{-}_2)^2}. \tag{J.48}$$

The two sectoral signals are then processed sequentially, the food-and-gas signal first and the services signal second, each as a scalar update that acts directly on the Cholesky factor. Observing $y_{1,t}$, which loads only on $z_{1,t}$, the innovation variance is $S_{1,t} = P^{-}_{11,t} + R_{1,t}$ and the mean-update gains are $k_{11} = P^{-}_{11,t}/S_{1,t}$ and $k_{21} = P^{-}_{12,t}/S_{1,t}$, so that with innovation $\nu_{1,t} = z_{1,t} - \hat{a}_{1,t} + \sqrt{R_{1,t}}\,\varepsilon^{sig1}_t$,

$$m^{(1)}_{1,t} = \hat{a}_{1,t} + k_{11}\nu_{1,t}, \qquad m^{(1)}_{2,t} = \hat{a}_{2,t} + k_{21}\nu_{1,t}. \tag{J.49}$$

Writing $P^{-}_t = \ell^{-}\ell^{-\top}$ with columns $c_1 = (\ell^{-}_1, \ell^{-}_2)^{\top}$ and $c_2 = (0, \ell^{-}_3)^{\top}$, and using $P^{-}_t e_1 = \ell^{-}_1 c_1$ for $e_1 = (1,0)^{\top}$, the scalar covariance update collapses the first column and leaves the second untouched,

$$P^{a}_t = P^{-}_t - \frac{(P^{-}_t e_1)(P^{-}_t e_1)^{\top}}{S_{1,t}} = \Big(1 - \tfrac{P^{-}_{11,t}}{S_{1,t}}\Big) c_1 c_1^{\top} + c_2 c_2^{\top} = \frac{R_{1,t}}{S_{1,t}}\, c_1 c_1^{\top} + c_2 c_2^{\top}, \tag{J.50}$$

because $1 - P^{-}_{11,t}/S_{1,t} = R_{1,t}/S_{1,t}$. The intermediate factor is therefore the prior factor with its first column rescaled by the square-root gain $g^{s}_{1,t} = \sqrt{R_{1,t}/S_{1,t}}$,

$$\ell^{a}_1 = g^{s}_{1,t}\, \ell^{-}_1, \qquad \ell^{a}_2 = g^{s}_{1,t}\, \ell^{-}_2, \qquad \ell^{a}_3 = \ell^{-}_3. \tag{J.51}$$

Observing $y_{2,t}$ next, which loads on $z_{2,t}$, the innovation variance is $S_{2,t} = (\ell^{a}_2)^2 + (\ell^{a}_3)^2 + R_{2,t}$, the gains are $k_{12} = \ell^{a}_1 \ell^{a}_2 / S_{2,t}$ and $k_{22} = \Big((\ell^{a}_2)^2 + (\ell^{a}_3)^2\Big)/S_{2,t}$, and with $\nu_{2,t} = z_{2,t} - m^{(1)}_{2,t} + \sqrt{R_{2,t}}\,\varepsilon^{sig2}_t$,

$$m_{1,t} = m^{(1)}_{1,t} + k_{12}\nu_{2,t}, \qquad m_{2,t} = m^{(1)}_{2,t} + k_{22}\nu_{2,t}. \tag{J.52}$$

The posterior Cholesky factor is available in closed form,

$$\ell^{+}_1 = \ell^{a}_1 \sqrt{\frac{(\ell^{a}_3)^2 + R_{2,t}}{S_{2,t}}}, \qquad \ell^{+}_2 = \frac{\ell^{a}_2\, R_{2,t}}{\sqrt{S_{2,t}\Big((\ell^{a}_3)^2 + R_{2,t}\Big)}}, \qquad \ell^{+}_3 = \ell^{a}_3 \sqrt{\frac{R_{2,t}}{(\ell^{a}_3)^2 + R_{2,t}}}, \tag{J.53}$$

and one verifies directly that $\ell^{+}\ell^{+\top}$ reproduces the Joseph-form posterior $\frac{P^{a}_t - (P^{a}_t e_2)(P^{a}_t e_2)^{\top}}{S_{2,t}}$

with $e_2 = (0,1)^\top$. The decisive feature of this form is that ℓ_1^+ and ℓ_3^+ are square roots of strictly positive quantities whenever $R_{2,t} > 0$, so the posterior factor has a strictly positive diagonal and $P_t = \ell^+ \ell^{+\top}$ is positive definite by construction at every point of the state space. A covariance-form filter would only be positive definite up to the perturbation error, and the residual indefiniteness feeds the third moment; the square-root form is what keeps the posterior positive-definite at every perturbation order; in earlier third-order experiments it is also what removed the order-three blow-up in the reported inflation-forecast variance. The common demand state is filtered as an independent scalar problem, with prior variance $P^-_{dd,t} = \rho^2_{zd} e^{2cLd_{t-1}} + \sigma^2_{zd}$, gain $k_d = P^-_{dd,t}/(P^-_{dd,t} + R_{d,t})$, innovation $\nu_{d,t} = z_{d,t} - \hat{a}_{d,t} + \sqrt{R_{d,t}}\,\varepsilon^{sigd}_t$ (where $R_{d,t} = \bar{R}_d \exp(q_{d,t})$ with $q_{d,t} = \sigma_{h,d}\,\varepsilon^h_{d,t}$ and $\sigma_{h,d} = 0$ at baseline, so the demand-signal noise is constant), posterior mean $m_{d,t} = \hat{a}_{d,t} + k_d \nu_{d,t}$, and posterior factor $\ell^+_d = \sqrt{P^-_{dd,t} R_{d,t}/(P^-_{dd,t} + R_{d,t})}$.

The updated factors are carried into the next period as the Cholesky states, with the diagonals stored in logs and the off-diagonal in levels,

$$cL1_t = \log \ell_1^+, \qquad cL2_t = \ell_2^+, \qquad cL3_t = \log \ell_3^+, \qquad cLd_t = \log \ell_d^+, \tag{J.54}$$

and the current posterior covariance is reconstructed from them exactly as its lagged counterpart was,

$$P_{11,t} = e^{2cL1_t}, \qquad P_{12,t} = e^{cL1_t} cL2_t, \qquad P_{22,t} = cL2_t^2 + e^{2cL3_t}, \qquad P_{dd,t} = e^{2cLd_t}. \tag{J.55}$$

These posterior objects map into the reported measures of uncertainty through the loadings of forecast inflation on the filtered states. Let $\mathbf{a}_j = (a_{j,m_1}, a_{j,m_2}, a_{j,m_d})^\top$ collect the sensitivities of the household's forecast of cumulative four-quarter sector-j inflation to the three posterior means $(m_{1,t}, m_{2,t}, m_{d,t})$, evaluated at the steady state. The posterior variance of that forecast is the quadratic form

$$sv_{j,t} = \mathbf{a}_j^\top P_t\, \mathbf{a}_j = a^2_{j,m_1} P_{11,t} + a^2_{j,m_2} P_{22,t} + 2\, a_{j,m_1} a_{j,m_2} P_{12,t} + a^2_{j,m_d} P_{dd,t}, \tag{J.56}$$

and the cross-sector posterior covariance of the two forecasts is the bilinear form $scov_{12,t} = \mathbf{a}_1^\top P_t\, \mathbf{a}_2$. Sectoral inflation uncertainty is the log posterior variance $h_{j,t} = \log sv_{j,t}$, with steady-state value $\bar{h}_j$ so that the policy-relevant gap entering the interest-rate rule is $h_{j,t} - \bar{h}_j$; the cross-sector uncertainty correlation is $\rho^{12}_t = scov_{12,t}/\sqrt{sv_{1,t}\, sv_{2,t}}$; and basket uncertainty, demeaned by its steady-state level $\bar{U}$, is

$$U_t = \omega_U^2\, sv_{1,t} + (1-\omega_U)^2\, sv_{2,t} + 2\,\omega_U(1-\omega_U)\, scov_{12,t} - \bar{U}, \tag{J.57}$$

where ω_U is the salience weight the basket places on food and gasoline; at baseline $\omega_U = \omega_1 = 0.27$, so this is the main text's Equation (17), and only the salience robustness of Appendix J.14 sets $\omega_U = 0.57$.

The signal-noise variance $R_{j,t}$ that drives this filter is not a hand-set precision loading but the trend-inflation price-dispersion object of Appendix J.3. There $R_{j,t} = \bar{R}_j \exp(q_{j,t})$, with $\bar{R}_j$ the steady-state signal-noise level and $q_{j,t}$ the log signal-noise state that tracks the year-on-year cross-item dispersion of sector-j price changes through the common pass-through κ_{disp} (with services signal noise additionally loading on the food-and-gas dispersion through κ_{serv}); the demand signal noise $R_{d,t} = \bar{R}_d \exp(q_{d,t})$ carries no policy

content. Because positive trend inflation $\bar{\pi}$ makes dispersion respond to sectoral inflation already at first order, with sensitivity $2\alpha_j\bar{\pi}/(1-\alpha_j)$ in the sector-j Calvo probability α_j (Appendix J.3), monetary policy moves posterior uncertainty along the composed chain from the policy stance to sectoral inflation, to price-change dispersion, to $q_{j,t}$ and hence $R_{j,t}$, and finally through the square-root update to P_t, $sv_{j,t}$, and $h_{j,t}$. The local policy semi-elasticity of sectoral uncertainty is therefore a solved equilibrium object of this chain rather than a free constant, and its cross-sector ordering is inherited from the primitive stickiness $\alpha_1 < \alpha_2$ acting through a single common pass-through κ.

J.6 Welfare from Primitives

The policy comparisons in Section 6.3 rank rules by the household's own lifetime utility, evaluated on the primitives of the model rather than through an augmented quadratic loss. Period felicity is

$$u_t = \xi_t \frac{C_t^{1-\sigma_c}}{1-\sigma_c} - \chi_t \frac{N_t^{1+\varphi}}{1+\varphi}, \tag{J.58}$$

where C_t is consumption, N_t hours, $\sigma_c = 1.5$ the inverse intertemporal elasticity of substitution, $\varphi = 2.0$ the inverse Frisch elasticity, χ_t the exogenous labour-disutility shifter, and ξ_t the effective marginal-utility-of-consumption shifter. The latter combines an exogenous preference innovation $\log \xi_t$ with the household's perception of the common demand state through a scalar loading eff^d, namely $\xi_t = \exp(\xi_t^v)$ with

$$\xi_t^v = \log \xi_t + \text{eff}^d \, d_t, \qquad d_t = \lambda^{hh} \, m_{d,t} + (1-\lambda^{hh}) \, z_{d,t}, \tag{J.59}$$

where d_t is the effective common-demand state the household acts on, $z_{d,t}$ is the latent common demand state, and $m_{d,t}$ is its posterior mean under the household's signal-extraction problem. At the baseline information share $\lambda^{hh} = 1$, so $d_t = m_{d,t}$: households evaluate their marginal utility at the posterior belief $m_{d,t}$ about the latent state $z_{d,t}$, not at $z_{d,t}$ itself.

Writing $c_t \equiv \log C_t$ and $n_t \equiv \log N_t$, the consumption and labour branches of lifetime utility satisfy the recursions

$$Wc_t = \frac{e^{\xi_t^v} \, e^{(1-\sigma_c)c_t}}{1-\sigma_c} + \beta \, \mathbb{E}_t \, Wc_{t+1}, \tag{J.60}$$

$$Wn_t = \frac{e^{\log \chi_t} \, e^{(1+\varphi)n_t}}{1+\varphi} + \beta \, \mathbb{E}_t \, Wn_{t+1}, \tag{J.61}$$

with discount factor $\beta = 0.99$, and perceived lifetime welfare is their difference,

$$W_t = Wc_t - Wn_t. \tag{J.62}$$

No uncertainty penalty is added by hand. Equations (J.60)–(J.62) contain only the household's own primitives, and every effect of sectoral inflation uncertainty on welfare reaches W_t endogenously through the equilibrium paths of consumption, hours, sectoral inflation, relative prices, and Calvo price dispersion. Evaluated at the deterministic steady state the two branches collapse to the closed-form perpetuities $\overline{Wc} = e^{\bar{\xi}} e^{(1-\sigma_c)c_{ss}}/[(1-\sigma_c)(1-\beta)]$ and $\overline{Wn} = e^{\bar{\chi}} e^{(1+\varphi)n_{ss}}/[(1+\varphi)(1-\beta)]$, which anchor the recursion.

Because households act on the belief $m_{d,t}$ rather than the realised state $z_{d,t}$, the perceived

object W_t in (J.62) values a given allocation at the household's own, generally mistaken, marginal utility. The object relevant for ranking policies is instead the true lifetime utility that the chosen allocation actually delivers. Let

$$\xi_t^{v,\text{true}} = \log \xi_t + \text{eff}^d \, z_{d,t} \tag{J.63}$$

be the marginal-utility shifter evaluated at the true demand state $z_{d,t}$, and define the true-state consumption branch

$$Wc_t^{\text{true}} = \frac{e^{\xi_t^{v,\text{true}}} \, e^{(1-\sigma_c)c_t}}{1-\sigma_c} + \beta \, \mathbb{E}_t \, Wc_{t+1}^{\text{true}}, \tag{J.64}$$

in which the allocation (c_t, n_t) is exactly the one households choose while acting on their noisy posterior beliefs $m_{d,t}$, whereas the felicity attached to that allocation is scored at the true state. Since hours enter no demand shifter and χ_t is exogenous, the labour branch Wn_t is common to both valuations, and the headline welfare object is

$$W_t^{\text{true}} = Wc_t^{\text{true}} - Wn_t. \tag{J.65}$$

This is the standard planner criterion under imperfect information: the true expected lifetime utility of an allocation that is itself chosen under noisy beliefs. The gap between W_t^{true} and the perceived W_t is precisely the households' misperception of the common demand state, entering through the term $\text{eff}^d(m_{d,t} - z_{d,t})$; the two objects coincide when beliefs are correct and diverge only through this demand-misperception term, so W_t^{true} charges welfare for belief errors that the household does not internalise. We report the perceived W_t alongside it as a cross-check.

Policies are ranked by expected true welfare $\mathbb{E}[W_t^{\text{true}}]$ expressed in consumption-equivalent units. For a policy-coefficient vector ϕ, let $\overline{W}^{\text{true}}(\phi) \equiv \mathbb{E}[W_t^{\text{true}}]$ denote its ergodic-mean true welfare and $\overline{Wc}^{\text{true}}(\phi)$, $\overline{Wn}(\phi)$ the corresponding sub-aggregates. The consumption-equivalent loss of a rule ϕ relative to the optimal rule $\phi^\star$ is the constant fraction $\lambda(\phi)$ of per-period consumption that, surrendered under $\phi^\star$ at every date and state, leaves the household indifferent to living under ϕ. A proportional transfer scales $C_t^{1-\sigma_c}$ by $(1-\lambda)^{1-\sigma_c}$ uniformly and leaves hours untouched, so $\lambda(\phi)$ solves

$$(1-\lambda)^{1-\sigma_c} \, \overline{Wc}^{\text{true}}(\phi^\star) - \overline{Wn}(\phi^\star) = \overline{W}^{\text{true}}(\phi), \tag{J.66}$$

which inverts in closed form to

$$\lambda(\phi) = 1 - \left(\frac{\overline{W}^{\text{true}}(\phi) + \overline{Wn}(\phi^\star)}{\overline{Wc}^{\text{true}}(\phi^\star)} \right)^{\frac{1}{1-\sigma_c}}. \tag{J.67}$$

On this metric a rule that fully looks through food/gas inflation uncertainty costs about 25 basis points of consumption-equivalent welfare relative to the optimum, and the conventional headline rule (expenditure-weighted inflation, no uncertainty response) about 35 basis points every period.

The recursions (J.60)–(J.65) are solved jointly with the rest of the model by pruned perturbation, and $\mathbb{E}[W_t^{\text{true}}]$ is computed as the ergodic mean of W_t^{true}; the policy evaluation is carried out at the second-order pruned solution. The ranking of rules does not

hinge on the approximation order, and the reason is structural rather than numerical. Uncertainty is not recovered as a perturbation residual but carried as an explicit belief-variance state: the sectoral inflation-forecast variances $\mathrm{sv}_{j,t} = a_j' P_t a_j$, with a_j the vector of cumulative-inflation loadings on the latent states and P_t the posterior covariance, and the basket-uncertainty gap U_t built from P_t. Fluctuations in these belief-variance states enter W_t^{true} through the equilibrium allocation at any approximation order, so the level of welfare, and hence the ranking, does not hinge on the higher-order stochastic corrections a pruned expansion adds. Because the posterior means $m_{1,t}, m_{2,t}, m_{d,t}$ and the Cholesky covariance states are themselves elements of the state vector, the order-3 solution loads on both the latent states $(z_{1,t}, z_{2,t}, z_{d,t})$ and their posterior moments, and W_t^{true} is a well-defined function of the simulated history under the model-consistent solution.

A second-order expansion of the primitive felicity (J.58) around the deterministic steady state reproduces the familiar log-linear (linear-quadratic) welfare-from-primitives representation, in which the period loss is quadratic in the output gap, sectoral inflation, and Calvo price dispersion, and the uncertainty wedges derived above, the Euler precautionary term and the recursive Calvo pricing-risk term, enter as the second moments that scale those quadratic pieces. That representation is useful as motivation, since it makes transparent why sectoral inflation uncertainty is costly and through which channels. The headline numbers, however, use the exact recursive object (J.65) evaluated on the nonlinear simulated paths rather than the quadratic proxy, so that no approximation is imposed on the felicity that is being compared across rules.

J.7 The Household and Firm Channels

Sectoral inflation uncertainty reaches welfare through two channels that are far from equal. On the household (demand) side, households observe the common demand state $z_{d,t}$ with noise and set consumption on their posterior mean $m_{d,t}$ through the Euler equation (J.15), so demand-belief errors translate into consumption that is misaligned with the true state. On the firm (supply) side, firms observe sectoral markup pressure $z_{j,t}$ with noise and set Calvo reset prices on their posterior means $m_{j,t}$ through the pricing block (J.21), so markup-belief errors translate into mispriced resets. The model isolates the two by letting each agent condition on the true state rather than on its posterior mean, which switches that agent's belief-error channel off.

Evaluated at the optimal rule and the estimated parameters, the welfare cost of noisy information relative to full information is 0.131% of consumption (about 13 basis points) in consumption-equivalent units, split 97 percent firm-side and 3 percent household-side; the decomposition is exactly additive across the two only-one-side-noisy configurations. Because firms price on the same noisy beliefs households hold, the burden of household inflation uncertainty is borne mainly through mispriced Calvo resets and the relative-price misallocation they create; households' own consumption-allocation errors carry the remainder. The split is sensitive to the signal-noise parameterization: the estimated measured-dispersion block routes gas–food relative-price volatility into household signal noise, which raises the household share relative to specifications without it, while per unit of belief noise the firm channel remains the costlier for three reasons. First, a firm's pricing error is locked in until its next Calvo adjustment, on average 3.3 quarters in food and gasoline and 12.5 in services, whereas households re-optimize consumption every period, so demand-belief errors are transitory and smoothed while pricing errors persist.

Second, mispriced resets create relative-price dispersion across varieties and sectors, the classic first-order welfare cost of a New Keynesian economy: dispersion raises the labor needed to produce a given consumption bundle, entering sectoral hours through the term $\log(1+\Delta_{j,t})$ in Equation (J.17), whereas a demand-belief error is a smoothable level effect on consumption. Third, the calibration places more variance on the markup states that firms must infer ($\sigma_z = 0.118$) than on the demand state households infer ($\sigma_{zd} = 0.040$). The inflation uncertainty the central bank leans against therefore combines firm-side pricing risk transmitted through relative-price misallocation with household-side belief noise driven by the salient gas price; neither is a precautionary-demand effect. That this is a genuine uncertainty channel, rather than a relabelling of some other friction, can be seen by scaling the signal-noise level: halving and doubling it moves the measured uncertainty and the firm-channel welfare cost in almost exactly the same proportion, so the cost is proportional to the posterior forecast variance. This is the exact sense in which the mechanism operates through uncertainty: for a Bayesian forecaster the posterior variance the survey measures is the expected squared forecast error, and it is those forecast errors, baked into sticky prices, that misallocate.

The decomposition holds the policy rule fixed at its optimum and measures the welfare cost of each agent's belief errors given that policy responds to the same measured uncertainty. It reinforces the case for leaning against uncertainty: the costs on both sides stem from the price-change dispersion that is the source of the agents' signal noise, which a tighter stance compresses (Appendix J.3); neither is a demand-side precautionary effect. It is also why the response is a pure risk-management instrument rather than a conventional inflation reaction: a shock to posterior uncertainty leaves the mean beliefs, and hence the first-order paths of sectoral inflation and the output gap, unchanged except through the policy rule's own reaction, so the impulse response of inflation and output to an uncertainty shock is identical whether the belief-error channels are on or off.

Third-order approximation and uncertainty wedges

We restate that the wedges in this subsection are a motivating analytical decomposition of the exact nonlinear Euler and Calvo blocks; they are not hand-added to the equations solved numerically in Section J.11. The second-order pruned perturbation of the live model produces the variance wedges endogenously, with loadings on the one-period conditional variance $\mathrm{Var}_t(\pi_{t+1})$ (and $\mathrm{Var}_t(\pi_{j,t+1})$ in the sectoral NKPC). We continue to write the loadings using U_t under the maintained closure $\mathrm{Var}_t(\pi_{t+1}) \approx U_t$ and flag this explicitly where it bites.

For the welfare analysis and for thinking about optimal policy, the third-order cumulant expansion of the equilibrium conditions makes the role of uncertainty explicit. The expanded IS equation takes the form

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma_c}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n\right) - \psi_U^{(2)} \mathrm{Var}_t(\pi_{t+1}) - \psi_{\mathrm{cov}}^{(2)} \Gamma_t - \psi_U^{(3)} S_{\pi,t} + \varepsilon_t^d, \quad \text{(J.68)}$$

$$\Gamma_t \equiv \sigma_c^2 \mathrm{Var}_t(\Delta c_{t+1}) - 2\sigma_c \mathrm{Cov}_t(\Delta c_{t+1}, \pi_{t+1}), \qquad S_{\pi,t} \equiv \kappa_3(\pi_{t+1} \mid t). \quad \text{(J.69)}$$

Here Γ_t collects second-moment cross terms between consumption growth and inflation, while $S_{\pi,t}$ captures inflation skewness. Solving for x_t from the log-Euler equation requires dividing by σ_c, so under a log-normal benchmark with small higher-order cumulants

the baseline coefficients satisfy $(\psi_U^{(2)}, \psi_{\text{cov}}^{(2)}, \psi_U^{(3)}) = (\frac{1}{2\sigma_c}, \frac{1}{2\sigma_c}, \frac{1}{6\sigma_c})$, so that higher inflation uncertainty and skewness both raise the risk premium and depress demand.

At the sectoral level, the New Keynesian Phillips curves generalize to (with $\sigma^2_{\pi,j,t}$ identified with $\text{Var}_t(\pi_{j,t+1})$ under the same maintained closure)

$$\pi_{j,t} = \beta \mathbb{E}_t \pi_{j,t+1} + \kappa_j \left(\phi_j x_t + \psi_j \tilde{p}_{j,t}\right) + \varphi_j^{(2)} \text{ Var}_t(\pi_{j,t+1}) + \varphi_j^{(3)} S_{\pi,j,t} + u_{j,t}, \tag{J.70}$$

$$\tilde{p}_{j,t} \equiv p_{j,t} - p_t, \qquad S_{\pi,j,t} \equiv \kappa_3(\pi_{j,t+1} \mid t), \qquad \kappa_j \equiv \frac{(1-\alpha_j)(1-\alpha_j\beta)}{\alpha_j}. \tag{J.71}$$

The coefficients $(\varphi_j^{(2)}, \varphi_j^{(3)})$ measure how sectoral inflation variance and skewness affect optimal markups. They are key for welfare, because they determine how strongly sector-specific uncertainty translates into inflation and relative-price distortions.

These wedges load on the one-period conditional variances $\text{Var}_t(\pi_{j,t+1})$, which the model identifies with the posterior inflation-forecast variances $\sigma^2_{\pi,j,t} = \text{sv}_{j,t} = \mathbf{a}_j' P_t \mathbf{a}_j$ produced by the square-root Kalman filter of Appendix J.5, with steady-state levels $\bar{\sigma}^2_{\pi,j} = e^{\bar{h}_j}$ anchoring the log-uncertainty gaps $\tilde{h}_{j,t} \equiv h_{j,t} - \bar{h}_j$. The policy sensitivity of these variances is not a free loading but the solved dispersion chain of Appendix J.3: positive trend inflation makes the year-on-year cross-item dispersion of price changes respond to sectoral inflation at first order, with sensitivity $2\alpha_j\bar{\pi}/(1-\alpha_j)$, so a tighter policy stance lowers sectoral inflation, compresses dispersion, and through the pass-through κ lowers $\sigma^2_{\pi,j,t}$. Expanding the level variance and the basket object around the steady state,

$$\sigma^2_{\pi,j,t} = e^{h_{j,t}} = \bar{\sigma}^2_{\pi,j}\left[1 + \tilde{h}_{j,t} + \tfrac{1}{2}\tilde{h}^2_{j,t} + \tfrac{1}{6}\tilde{h}^3_{j,t} + O(4)\right], \tag{J.72}$$

$$U_t = \omega_1^2 \sigma^2_{\pi,1,t} + (1-\omega_1)^2 \sigma^2_{\pi,2,t} + 2\omega_1(1-\omega_1)\rho_t \sqrt{\sigma^2_{\pi,1,t}\sigma^2_{\pi,2,t}}, \tag{J.73}$$

maps these belief-variance fluctuations, driven by the dispersion channel, into the IS and Phillips-curve wedges of Equations (J.68)–(J.70) that enter the welfare-cost calculations.

For completeness, the third-order IS wedge can be derived explicitly from the nonlinear Euler equation. Let

$$Z_{t+1} \equiv \log(\xi_{t+1}/\xi_t) - \sigma_c \Delta c_{t+1} - \pi_{t+1}, \qquad \Delta c_{t+1} \equiv \log C_{t+1} - \log C_t.$$

Taking logs of the Euler condition and applying a cumulant expansion up to third order gives

$$0 = \log\beta + i_t - \tfrac{1}{2}i_t^2 + \tfrac{1}{3}i_t^3 + \mathbb{E}_t[Z_{t+1}] + \tfrac{1}{2}\text{Var}_t(Z_{t+1}) + \tfrac{1}{6}\kappa_3(Z_{t+1}) + O(4), \tag{J.74}$$

$$\text{Var}_t(Z_{t+1}) = \sigma_c^2 \text{Var}_t(\Delta c_{t+1}) + \text{Var}_t(\pi_{t+1}) - 2\sigma_c \text{Cov}_t(\Delta c_{t+1}, \pi_{t+1}), \tag{J.75}$$

$$\kappa_3(Z_{t+1}) = \kappa_3\left(\log \tfrac{\xi_{t+1}}{\xi_t}\right) - \sigma_c^3 \kappa_3(\Delta c_{t+1}) - \kappa_3(\pi_{t+1}) \tag{J.76}$$

$$+ 3\sigma_c^2 \text{Cum}(\Delta c_{t+1}, \Delta c_{t+1}, \pi_{t+1}) - 3\sigma_c \text{Cum}(\Delta c_{t+1}, \pi_{t+1}, \pi_{t+1}). \tag{J.77}$$

Mapping $\mathbb{E}_t \Delta c_{t+1}$ to $\mathbb{E}_t x_{t+1} - x_t$ and collecting terms yields the coefficients $(\psi_U^{(2)}, \psi_{\text{cov}}^{(2)}, \psi_U^{(3)})$ in (J.68). These coefficients govern the marginal welfare cost of increases in inflation variance, covariance terms, and skewness.

Parameterization and calibration

Table J.8 summarizes the model's baseline calibration. Preference and technology parameters take standard quarterly New Keynesian values. The discount factor $\beta = 0.99$ implies a four percent annualized real rate (Galí, 2015); risk aversion $\sigma_c = 1.5$ and the inverse Frisch elasticity $\varphi = 2.0$ follow Smets and Wouters (2007); and the within-sector elasticity of substitution $\theta = 6$ implies a twenty percent steady-state markup, as in Christiano et al. (2005) and Galí (2015). The elasticity of substitution across the two sectoral bundles is $\eta = 2.0$. The expenditure weight $\omega_1 = 0.27$ is derived from SCE category weights (food 17.83% plus transportation 10.09%), normalized by the sum of all reported category means (103.90%, which exceeds 100% because respondents report each category independently and reported shares need not sum to one); this yields $\omega_1 = (17.83 + 10.09)/103.90 \approx 0.269$. Positive trend inflation $\bar{\pi} = \log(1.02)/4 \approx 0.00495$ is what activates the price-dispersion channel at first order, so it is held fixed at a two-percent annual target rather than calibrated.

The sectoral Calvo non-adjustment probabilities $(\alpha_1, \alpha_2) = (0.70, 0.92)$ are set so that the model reproduces the realized sectoral inflation-volatility ratio in the BLS CPI data, in which food and gas prices are far more flexible than services: food and gas (Sector 1) reset on average every ≈ 3.3 quarters and services (Sector 2) every ≈ 12.5 quarters, bracketing the regular-price evidence in Nakamura and Steinsson (2008). Because the first-order loading of year-on-year price-change dispersion on sectoral inflation is $2\alpha_j\bar{\pi}/(1-\alpha_j)$, these two rigidities also fix the relative strength of the dispersion channel across sectors, at 0.023 for food and gas and 0.114 for services.

The central transmission parameter is the dispersion pass-through κ_{disp}, which maps proportional deviations of sectoral price-change dispersion into the log signal-noise variance $q_{j,t}$ and thus into posterior inflation-forecast uncertainty $h_{j,t}$. It is estimated at $\kappa_{\text{disp}} = 2.08$ jointly with the salience loading $\kappa_{\text{serv}} = 0.83$, the memory ρ_q, and the measured-dispersion block by minimum distance (Appendix J.13); at the estimate the cumulative food-and-gas uncertainty response is -9.9 against -9.2 in the data, within the 90% band at every quarter. The salience loading κ_{serv} fits the services trough size, while its timing, imported from the food-and-gas dispersion dynamics, is a prediction: the model produces a trough of -1.9% in the third quarter, inside the wide empirical band around the -2.3% trough-quarter point estimate, with a services-to-food-and-gas response ratio of about 0.42 against 0.40 in the data, implied jointly by the two fitted responses rather than separately targeted. Because food and gas inflation moves far more under policy than services inflation does (its Phillips curve is steep), the net effect running from policy through dispersion to uncertainty is larger for food and gas, and this asymmetry is inherited entirely from (α_1, α_2) rather than imposed.

The latent markup states are correlated across sectors, $z_{2,t} = \rho_z z_{2,t-1} + \sigma_z(\text{corr}_z\, \varepsilon_t^{z_1} + \sqrt{1-\text{corr}_z^2}\, \varepsilon_t^{z_2})$, and this is the model's only cross-sector coupling. The coupling is measured against one moment: $\text{corr}_z = 0.055$ reproduces the near-zero correlation of realized sectoral inflation (0.04 in the data, 0.04 in the model). The within-household correlation of the two sectors' uncertainty measures, 0.34 in the SCE (0.26 averaging all category pairs), is carried separately, by the salience channel together with the idiosyncratic services-uncertainty scale $\sigma_{h,2}$, which is set so the model reproduces the 0.34 exactly; belief correlation and fundamental correlation are thus disciplined by different moments, as they are different objects. The markup and demand states have persistences $\rho_z = 0.55$

and $\rho_{zd} = 0.60$ with innovation standard deviations $\sigma_z = 0.118$ and $\sigma_{zd} = 0.040$, chosen so that realized sectoral inflation volatility is of the empirically relevant magnitude (Table 6 reports the resulting fit); the latent-state persistences ρ_z are fixed a priori at moderate quarterly values and matter mainly through the steady-state Kalman gains, while the conditional dynamics the estimation targets are carried by the estimated $(\rho_q, \kappa_{\text{disp}}, \kappa_{\text{serv}})$; the steady-state signal-noise variances are $\bar{R}_1 = \bar{R}_2 = 0.02$ and $\bar{R}_d = 0.002$, levels that have no separate data anchor: only proportional deviations of signal noise are matched by the estimation, and the homogeneous level rescaling of Appendix J.14 leaves the fit exactly invariant while scaling the welfare magnitudes; and the preference, labour, and technology shocks are switched off in the baseline so that the reported dynamics isolate the markup, demand, dispersion, and monetary-policy channels. The uncertainty wedge coefficients reported in the table remain useful as a log-linear interpretation of the mechanism, but the numerical solution is performed on the exact nonlinear Euler, Calvo, and Kalman blocks and does not use them.

J.8 Derivation of the Sectoral NKPC with Uncertainty Wedges

This section makes explicit how the NKPC uncertainty wedges arise from the Calvo reset-price condition and how they map to deep parameters $(\beta, \alpha_j, \theta_j)$.

Step 1: Express the reset price in real terms and isolate cumulative inflation. Fix a sector and suppress the j index. Let within-sector elasticity be $\theta > 1$ and Calvo non-adjustment probability be $\alpha \in (0, 1)$. Define the cumulative (log) sectoral inflation $\Pi_{t,t+s} \equiv \sum_{u=1}^{s} \pi_{t+u}$. Using the closed-form reset price and dividing by P_t, the reset price can be written as

$$\frac{p_t^*}{P_t} = \mu \cdot \frac{\sum_{s\geq 0}(\alpha\beta)^s \, \mathbb{E}_t[\Lambda_{t,t+s}\, mc_{t+s}\, Y_{t+s}\, \exp(\theta\, \Pi_{t,t+s})]}{\sum_{s\geq 0}(\alpha\beta)^s \, \mathbb{E}_t[\Lambda_{t,t+s}\, Y_{t+s}\, \exp((\theta-1)\, \Pi_{t,t+s})]}, \qquad \mu = \frac{\theta}{\theta - 1}. \tag{J.78}$$

Step 2: Cumulant expansion and the pure inflation-risk coefficients. For any random variable Z, a third-order cumulant expansion gives

$$\log \mathbb{E}_t[e^Z] = \mathbb{E}_t[Z] + \tfrac{1}{2} \operatorname{Var}_t(Z) + \tfrac{1}{6}\kappa_3(Z \mid t) + O(4).$$

In (J.78), the numerator and denominator differ in the coefficient on $\Pi_{t,t+s}$: θ vs. $\theta - 1$. Focusing on the pure inflation-risk contributions (abstracting from cross-cumulants with the SDF/demand/marginal cost), the difference in variance terms contributes

$$\frac{1}{2}\left(\theta^2 - (\theta-1)^2\right) \operatorname{Var}_t(\Pi_{t,t+s}) = \frac{2\theta - 1}{2} \operatorname{Var}_t(\Pi_{t,t+s}),$$

and the difference in third-cumulant terms contributes

$$\frac{1}{6}\left(\theta^3 - (\theta-1)^3\right) \kappa_3(\Pi_{t,t+s} \mid t) = \frac{3\theta^2 - 3\theta + 1}{6} \kappa_3(\Pi_{t,t+s} \mid t).$$

Thus the deep-parameter risk-loading polynomials are $\frac{2\theta-1}{2}$ (variance) and $\frac{3\theta^2-3\theta+1}{6}$ (skewness).

Table J.8: Baseline Parameter Calibration

Symbol	Value	Description
Structural Parameters		
β	0.99	Discount factor
σ_c	1.5	Inverse intertemporal elasticity of substitution
φ	2.0	Inverse Frisch elasticity of labour supply
η	2.0	Elasticity of substitution across the two sectoral bundles
ω_1	0.27	Expenditure share on good 1 (food/gas)
θ, θ_j	6.0	Elasticity of substitution across varieties
μ_j	1.20	Desired static markup $\theta_j/(\theta_j - 1)$
α_1, α_2	0.70, 0.92	Calvo price non-adjustment, sectors 1 (food/gas) and 2 (services)
κ_1, κ_2	0.132, 0.008	NKPC slope $(1-\alpha_j)(1-\alpha_j\beta)/\alpha_j$
$\bar{\pi}$	0.00495	Trend inflation $\bar{\pi} = \log(1.02)/4$ (2% annual target; activates the dispersion channel at first order)
Dispersion and Signal-Extraction		
κ_{disp}	2.08	Dispersion pass-through: loading of own-sector year-on-year price-change dispersion into the log signal-noise state $q_{j,t}$ (estimated by minimum distance, Appendix J.13)
κ_{serv}	0.83	Salience loading: food-and-gas dispersion in services signal noise (DAcunto et al., 2021) (estimated; identified by the services uncertainty response)
$\bar{R}_1, \bar{R}_2$	0.02	Steady-state signal-noise (observation) variance, sectors 1 and 2
$\bar{R}_d$	0.002	Steady-state signal-noise variance, common demand state
ρ_q	0.55	Persistence of the signal-noise state $q_{j,t}$ (estimated; measured-uncertainty persistence 0.89 against 0.88)
$\sigma_{h,1},\ \sigma_{h,2}$	0, 5.25	Idiosyncratic household-uncertainty innovations: food-and-gas off; services set so the within-household cross-sector uncertainty correlation matches the SCE (0.34)
corr_z	0.055	Cross-sector markup-innovation coupling (measured: reproduces the near-zero realized cross-sector inflation correlation, 0.04)
σ_{D1}, σ_{D2}	$1.45\times10^{-4},\ 9.99\times10^{-4}$	Std. devs. of the independent dispersion shocks (half of $\bar{D}_j^{YoY}$; move uncertainty orthogonally to the price level)
$\rho_d,\ \zeta_g,\ \sigma_g$	0.20, 0, 0.125	Gas–food relative-price block: persistence and oil-innovation volatility (estimated; activity share estimated at zero, unconditional gap s.d. 18pp imposed; Appendix J.4)
ψ_v	357	Policy sensitivity of oil-innovation volatility (interior estimate; between component's impact compression)
$b_{g2},\ \rho_v$	0.0068, 0.95	Between-term weight (measured between/within variance ratio 0.757) and slow normalization of the between variance
$\rho_b,\ \nu_d,\ \nu_c$	0.58, 0.013, 0.012	Deferred-adjustment stock and deferral/catch-up loadings in measured dispersion (estimated, catch-up share at its intensity cap; sign change of the dispersion response)
$\bar{h}_1, \bar{h}_2$	−10.44, −13.73	Steady-state log inflation-forecast variance (belief anchor implied by the SS Kalman covariance; subjective s.d. 0.54% / 0.10%)
$\bar{\rho}_{12}$	0.55	Steady-state cross-sector belief correlation (implied)
Policy Rule		
ρ_i	0.3	Interest-rate smoothing
ϕ_π	2.0	Response to the sectoral-weighted inflation gap
ϕ_x	0.0	Response to the output gap (no gap response in baseline)
σ_i	0.0025	Std. dev. of the monetary-policy shock
θ_1^*	0.0	Optimal level inflation weight (0 = look through food/gas prices, i.e. core targeting)
ϕ_{U1}^*, ϕ_{U2}^*	0.003, 0.0	Optimal sectoral uncertainty-response margins (lean against food/gas uncertainty only; the Taylor-only baseline sets both to 0)
Exogenous Processes		
ρ_z, σ_z	0.55, 0.118	Latent sectoral markup-pressure process
ρ_{zd}, σ_{zd}	0.60, 0.040	Common demand-pressure process
ρ_ξ, σ_ξ	0.7, 0.0	Intertemporal preference shock (off in baseline)
ρ_χ, σ_χ	0.7, 0.0	Labour-disutility shock (off in baseline)
$\rho_{A,j}, \sigma_{A,j}$	0.8, 0.0	Sectoral technology shocks (off in baseline)
Log-Linear Uncertainty Wedges *(interpretive; produced endogenously by the order-3 solve, not calibrated inputs)*		
$\varphi_j^{(2)}$	1.63, 0.44	NKPC variance wedge $\kappa_j \frac{\alpha_j\beta}{1-\alpha_j\beta}\frac{2\theta-1}{2}$ (sector 1, sector 2)
$\varphi_j^{(3)}$	4.51, 1.20	NKPC skewness wedge $\kappa_j \frac{\alpha_j\beta}{1-\alpha_j\beta}\frac{3\theta^2-3\theta+1}{6}$ (sector 1, sector 2)
$\psi_U^{(2)}$	0.33	IS uncertainty wedge load $= 1/(2\sigma_c)$
$\psi_{\text{cov}}^{(2)}$	0.33	IS covariance wedge load $= 1/(2\sigma_c)$
$\psi_U^{(3)}$	0.11	IS skewness wedge load $= 1/(6\sigma_c)$
ϕ_j, ψ_j	1.0, 0.0	Mapping of x_t and $\tilde{p}_{j,t}$ into $mc_{j,t}$

Step 3: Collapse horizon-s cumulants to one-step cumulants. To express wedges using one-step-ahead objects $\sigma^2_{\pi,t} \equiv \text{Var}_t(\pi_{t+1})$ and $S_{\pi,t} \equiv \kappa_3(\pi_{t+1} \mid t)$, we impose a closure where inflation innovations are approximately independent across horizons so

that cumulants add:

$$\mathrm{Var}_t(\Pi_{t,t+s}) \approx s\,\sigma^2_{\pi,t}, \qquad \kappa_3(\Pi_{t,t+s} \mid t) \approx s\,S_{\pi,t}.$$

Then the Calvo-duration weighting implied by linearizing the log-sum around steady state yields $(1-\alpha\beta)\sum_{s\geq 1}(\alpha\beta)^s s = \frac{\alpha\beta}{1-\alpha\beta}$. Hence the risk adjustment to the reset price becomes proportional to

$$\frac{\alpha\beta}{1-\alpha\beta}\left[\frac{2\theta-1}{2}\sigma^2_{\pi,t} + \frac{3\theta^2-3\theta+1}{6}S_{\pi,t}\right].$$

Step 4: From reset-price risk adjustment to the NKPC wedge. Log-linearizing the Calvo price index around zero inflation delivers the NKPC slope $\kappa = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}$. Therefore the implied NKPC wedge coefficients are:

$$\varphi^{(2)} = \kappa\frac{\alpha\beta}{1-\alpha\beta}\cdot\frac{2\theta-1}{2}, \qquad \varphi^{(3)} = \kappa\frac{\alpha\beta}{1-\alpha\beta}\cdot\frac{3\theta^2-3\theta+1}{6}. \tag{J.79}$$

J.9 Third-Order Derivation of the IS Uncertainty Wedge

Let $Z_{t+1} \equiv \log(\xi_{t+1}/\xi_t) - \sigma_c\Delta c_{t+1} - \pi_{t+1}$. Taking logs in the Euler equation and using the cumulant expansion up to third order:

$$0 = \log\beta + i_t - \tfrac{1}{2}i_t^2 + \tfrac{1}{3}i_t^3 + \mathbb{E}_t[Z_{t+1}] + \tfrac{1}{2}\mathrm{Var}_t(Z_{t+1}) + \tfrac{1}{6}\kappa_3(Z_{t+1}) + O(4). \tag{J.80}$$

Subtract the efficient counterpart and map $\mathbb{E}_t\Delta c_{t+1}$ to $\mathbb{E}_t x_{t+1} - x_t$ to obtain

$$\begin{aligned} x_t = \mathbb{E}_t x_{t+1} &- \frac{1}{\sigma_c}\left(i_t - \mathbb{E}_t\pi_{t+1} - r_t^n\right) - \tfrac{1}{2}\mathrm{Var}_t(\pi_{t+1}) \\ &- \tfrac{1}{2}\Big(\sigma_c^2\,\mathrm{Var}_t(\Delta c_{t+1}) - 2\sigma_c\,\mathrm{Cov}_t(\Delta c_{t+1},\pi_{t+1})\Big) \\ &- \tfrac{1}{6}\kappa_3(\pi_{t+1}) + (\text{higher-order cross-cumulants}) + \varepsilon_t^d, \end{aligned} \tag{J.81}$$

which summarizes to

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma_c}(i_t - \mathbb{E}_t\pi_{t+1} - r_t^n) - \psi_U^{(2)}U_t - \psi_{\mathrm{cov}}^{(2)}\Gamma_t - \psi_U^{(3)}S_{\pi,t} + \varepsilon_t^d, \tag{J.82}$$

with baseline coefficients $(\psi_U^{(2)}, \psi_{\mathrm{cov}}^{(2)}, \psi_U^{(3)}) = (\frac{1}{2\sigma_c}, \frac{1}{2\sigma_c}, \frac{1}{6\sigma_c})$ under log-normality. The cumulant expansion of the exact Euler delivers a Jensen correction proportional to $\mathrm{Var}_t(\pi_{t+1})$, the one-period-ahead conditional variance. We then write the IS wedge in terms of basket uncertainty U_t under the maintained closure $\mathrm{Var}_t(\pi_{t+1}) \approx U_t$, where U_t is the four-quarter posterior basket-variance object of Equation (17) of the main text. This closure is inherited from the reduced-form benchmark of the reduced-form log-linear benchmark and is the only sense in which the U_t that enters the Taylor rule (18) coincides with the U_t written in the wedge equations. In the live numerical solution of Section J.11 no such substitution is made: the pruned perturbation of the exact nonlinear Euler produces the Jensen term in $\mathrm{Var}_t(\pi_{t+1})$ directly, and posterior uncertainty enters the model only through the sectoral gaps $h_{j,t}$ in the Taylor rule (18).

J.10 Discussion of moment fit

Table 6 reports the moments entering the minimum-distance estimation and the moments the model matches. The dispersion pass-through is estimated at $\kappa_{\text{disp}} = 2.08$ against the household-survey impulse response of food-and-gasoline inflation uncertainty to a contractionary monetary policy shock (Section 5), matched over the three quarters the projections span because the mean-gap state distributes the response over several quarters: the model's cumulative response is -9.9 against a covariance-aware empirical -9.2 (relative-percent), it lies within the empirical 90% band at every quarter, and it builds gradually to its trough rather than jumping at impact, as in the data. The sectoral Calvo parameters $(\alpha_1, \alpha_2) = (0.70, 0.92)$ are set to the ratio of realized sectoral inflation volatilities. Two correlations are matched by measured constants: the markup coupling $\text{corr}_z = 0.055$ reproduces the near-zero correlation of realized sectoral inflation (0.04), and the idiosyncratic services-uncertainty scale $\sigma_{h,2}$ reproduces the within-household cross-sector correlation of the uncertainty measures (0.34 in the SCE, exactly matched). The size of the services uncertainty response is fitted through the salience loading $\kappa_{\text{serv}} = 0.83$; its timing is not, and the model's trough arrives in the third quarter as in the data, at -1.9 against an empirical trough-quarter response of -2.3 whose 90% confidence interval $[-10.0, 5.4]$ includes zero, with a trough ratio of 0.42 against 0.40. The remaining unconditional moments are reported as diagnostics rather than calibration targets.

Two features of the fit deserve comment. Under a rule that does not respond to uncertainty, the realized four-quarter sectoral inflation volatilities are 3.0 percentage points in food and gasoline, against a measured composite of 4.8, and 0.28 points in services, against 0.8; the volatility ratio that identifies (α_1, α_2) is 10.8 under the exact dispersion recursion, above the measured 6.1, and the model undershoots the levels of realized sectoral volatility, reflecting the common markup-shock scale inherited across sectors. The correlation of four-quarter cumulative sectoral inflation realizations is now matched by measurement, $\text{corr}(\sum_{k=1}^{4} \pi_{1,t+k}, \sum_{k=1}^{4} \pi_{2,t+k}) = 0.04$ in model and data, and the belief correlation is carried separately by the salience channel and the idiosyncratic services noise (0.34, matched); realized and belief correlation are different objects and are disciplined by different constants. Together with the smooth recovery of the modelled uncertainty responses, which does not reproduce the statistically insignificant rebound in the empirical food-and-gasoline response at the three-quarter horizon, a persistence feature orthogonal to the estimation, it accounts for the moments the model does not match. The steady-state level of subjective uncertainty (a posterior standard deviation of 0.54 and 0.10 percent for food-and-gas and services) is well below the average SCE density-forecast standard deviation of roughly 2.5 percent, which bundles idiosyncratic heterogeneity and reporting behaviour that the representative filter does not model; the calibration is insulated from this level gap because only proportional deviations are matched and κ_{disp} absorbs the units mapping (Section 6.2 of the main text). The realized-volatility level gap is likewise a scale rather than a shape issue: Appendix J.14 shows that the homogeneous volatility rescaling that closes it leaves the estimation exactly invariant and scales the welfare magnitudes about threefold. Consistent with the mechanism operating through accumulated realized dispersion rather than through the inflation level, the ergodic correlation between the log posterior variance of food-and-gasoline inflation $h_{1,t}$ and food-and-gasoline inflation $\pi_{1,t}$ is $\text{corr}(h_{1,t}, \pi_{1,t}) \approx -0.11$ under a rule that does not respond to uncertainty, so posterior uncertainty moves essentially orthogonally to the inflation rate.

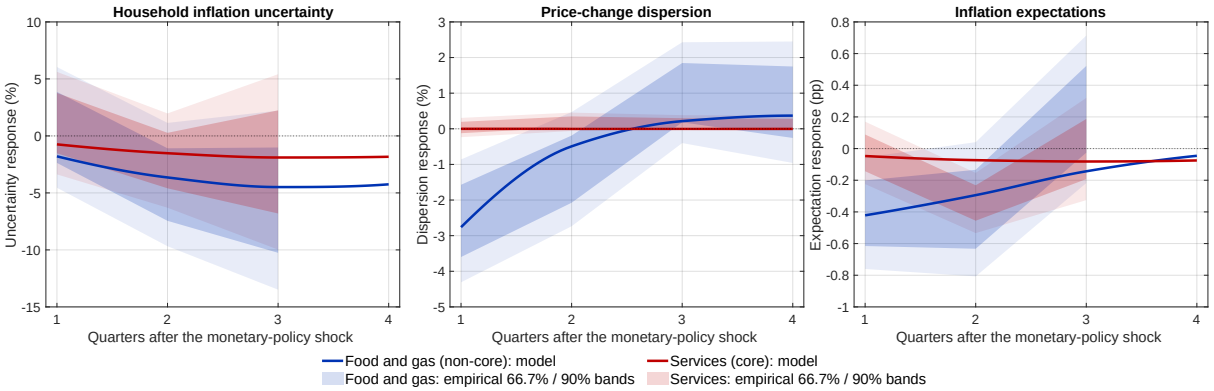


Figure J.14: Model versus Empirical Responses to a Contractionary Monetary Policy Shock: Uncertainty, Price-Change Dispersion, and Inflation Expectations

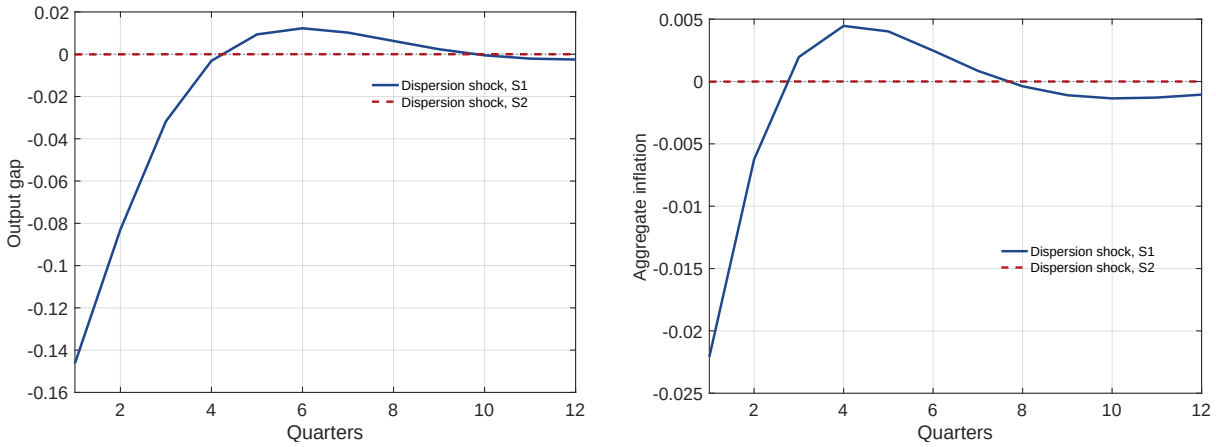
Note: Each panel overlays both sectors on a shared vertical scale: food and gas (non-core) in blue, services (core) in red. Shaded regions are the empirical 66.7% (dark) and 90% (light) local-projection bands; the central estimate is omitted so the model line reads directly against the bands. Solid coloured lines are the model responses. *Left (uncertainty):* the model's response of posterior four-quarter sectoral inflation variance at the estimated transmission parameters, shown in the same relative-percent units as the estimation target with no re-scaling: the food-and-gas response lies within the empirical 90% band at every quarter, and the services trough size, fitted through the salience loading κ_{serv} with its timing an unparameterized prediction, is -1.9 against an empirical -2.3. The empirical bands are the same covariance-aware respondent-level expenditure-weighted sector composites the model is estimated on (Section 5, the target of Section 6.2), averaged from months to quarters; the survey projections span three quarters and the model line is drawn to the fourth, and pooling to quarters dilutes the monthly significance, which is concentrated at the month-five trough. *Middle (dispersion):* the model's year-on-year price-change dispersion response, scaled to the salient (food-and-gas) impact quarter, against the BLS item-level dispersion local projections of Section 5.2 (salient = food, gas, energy; less salient = rent, medical, education); non-core dispersion compresses strongly while core dispersion barely moves. *Right (expectations):* the model's sectoral inflation path through the underreaction measurement equation of Coibion and Gorodnichenko (2015a) at $\lambda = 0.55$ plus the measured uncertainty premium of reported forecasts, against the SCE symmetric expectation local projections on the same composites. Scaling, bands, and the COVID treatment are as in Figure 11.

J.11 Simulation Results: Impact of Sectoral Uncertainty Shocks

To verify that the calibrated model behaves as the higher-order derivations predict, we simulate impulse responses from the pruned solution to a one-standard-deviation independent sectoral dispersion shock $\varepsilon^{D}_{j,t}$. This shock raises the year-on-year price-change dispersion $D^{YoY}_{j,t}$ in sector j (Equation (15)) without shifting the sectoral inflation mean, so it moves uncertainty orthogonally to the price level. Through the signal-noise mapping (Equation (16)) the extra dispersion raises the sector-j signal-noise state $q_{j,t}$ by the pass-through κ, raises the observation-noise variance $R_{j,t} = \bar{R}_j \exp(q_{j,t})$, and, through the Kalman update, raises the posterior conditional variance of sector-j inflation $\sigma^2_{\pi,j,t}$. The calibration follows the baseline values in Table J.8.

The simulated responses confirm the three economic channels emphasized in Section 6.1. First, the shock raises the year-on-year price-change dispersion $D^{YoY}_{j,t}$, which the signal-noise state tracks with its estimated memory ($\rho_q = 0.55$), so the observation-noise variance and, through the Kalman update, the posterior log-variance $h_{j,t} = \log \sigma^2_{\pi,j,t}$ rise on impact and decay with the dispersion itself and the state's own memory, the dispersion's persistence coming from the sector's Calvo variance recursion and the four-quarter year-on-year window; this mechanically raises aggregate basket uncertainty via

the expenditure-share-weighted aggregator. Second, the rise in the shocked sector's uncertainty gap $h_{j,t}$ tightens the Taylor-rule policy stance through the sectoral loading $\phi_{U,j}$; the resulting expected real-rate path depresses the output gap x_t on impact through the exact nonlinear Euler equation, consistent with a demand-channel impact elasticity of $-(1-\rho_i)\phi_{U,j}/\sigma_c$. The second-cumulant Jensen correction produced endogenously by pruning of the exact Euler is a separate, quantitatively smaller effect proportional to $\mathrm{Var}_t(\pi_{t+1})$. Third, higher posterior uncertainty raises the risk-adjusted desired markup inside the recursive Calvo block via the analogous second-cumulant Jensen correction on $E_t[\exp(\theta_j \pi_{j,t+1})]$, captured by the pruned perturbation, and propagates as additional inflationary pressure in the shocked sector. In the baseline calibration the policy-induced demand contraction dominates, so the aggregate inflation response π_t is small and slightly negative: an uncertainty shock that leaves the sectoral inflation mean untouched still lowers inflation on impact because policy leans against the higher posterior variance. This is why the uncertainty response in the policy rule is interpreted as risk management rather than as a conventional response to contemporaneous inflation pressure: policy leans against a state variable that raises posterior variance and future welfare losses, even when its short-run inflation effect is muted or negative. The non-shocked sector responds primarily to the endogenous fall in the output gap, while the shocked sector also inherits the direct markup-wedge contribution, so the cross-sector contrast in $\pi_{j,t}$ traces out the sectoral pricing-risk amplifier that drives the asymmetric optimal rule of Section 6.3.



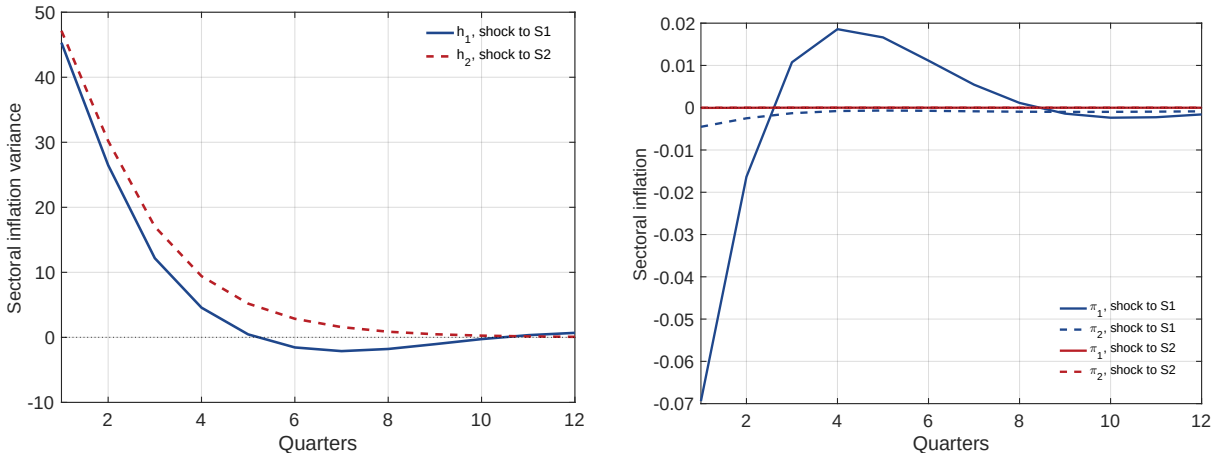
(a) Output gap (b) Aggregate inflation

Figure J.15: Model IRFs to Sectoral Dispersion Shocks: Aggregate Variables

Note: Responses are generated from the pruned solution of the two-sector model. Blue lines correspond to a one-standard-deviation independent dispersion shock $\varepsilon^D_{1,t}$ in Sector 1; red dashed lines correspond to the same shock $\varepsilon^D_{2,t}$ in Sector 2. The shock raises year-on-year price-change dispersion $D^{YoY}_{j,t}$ and hence posterior inflation uncertainty without moving the sectoral inflation mean.

J.12 Robustness of the Inversion to the Expenditure Share

The baseline optimal-policy analysis of Section 6.3 uses the empirical expenditure share $\omega_1 = 0.27$ and the baseline sectoral Calvo rigidities $(\alpha_1, \alpha_2) = (0.70, 0.92)$. Because the food-and-gas budget share is small, one might worry that the inversion, look through the food-and-gas price level and lean on its uncertainty, is an artifact of that weighting. To check this we re-optimize the two sectoral uncertainty responses $(\phi_{U,1}, \phi_{U,2})$ under



(a) Sectoral inflation variance

(b) Sectoral inflation

Figure J.16: Model IRFs to Sectoral Dispersion Shocks: Sectoral Variables

Note: Responses are generated from the pruned solution of the two-sector model. Solid lines plot Sector 1 responses and dashed lines plot Sector 2 responses. Blue lines correspond to a Sector 1 dispersion shock $\varepsilon^D_{1,t}$; red lines correspond to a Sector 2 dispersion shock $\varepsilon^D_{2,t}$.

equal expenditure shares $\omega_1 = 0.50$, holding the sectoral rigidities and everything else at baseline and targeting core inflation throughout ($\theta_1 = 0$).

Table J.9 reports the welfare-optimal sectoral uncertainty responses. In every configuration the optimum leans on food-and-gas uncertainty and looks through services uncertainty, $\phi^*_{U,1} > 0 = \phi^*_{U,2}$, and doubling the expenditure share moves the optimal food-and-gas response only from 0.003 to 0.005: a positive response to food-and-gas uncertainty together with a zero response to services uncertainty is optimal whether food and gas make up a quarter or a half of the basket. A full look-through of food-and-gas uncertainty ($\phi_{U,1} = 0$) remains sharply welfare-dominated at every share. The result does not rest on the expenditure weighting but on the flexibility gap $\alpha_1 < \alpha_2$ (Appendix J.3): monetary policy compresses price-change dispersion, and hence the uncertainty households attach to a sector, far more in flexible food and gas than in sticky services, and food-and-gas inflation uncertainty is the more variable object to begin with ($\bar{h}_1 \gg \bar{h}_2$). Leaning on services uncertainty buys negligible stabilization because policy can barely move it, so the optimal services response stays at zero regardless of the sector's budget weight.

J.13 Estimation: Moments, Identification, and Numerical Implementation

The transmission parameters $\theta = (\kappa_{\text{disp}}, \kappa_{\text{serv}}, \rho_q, \rho_d, \psi_v, \rho_b, \nu_c/\nu_d, \nu_d, \text{share}_x)$ are estimated by minimum distance, $\hat{\theta} = \arg\min [m(\theta) - \hat{m}]' W [m(\theta) - \hat{m}]$ with W diagonal at the inverse squared standard errors. The moment vector stacks: the food-and-gas and services uncertainty impulse responses, quarters one to three, in relative-percent units so κ_{disp} is identified (the empirical composites embed within-respondent covariances: the expenditure-weighted sector composite is built at the respondent level, the local projection of Section 5 is run on the composite, and monthly coefficients are averaged within quarters); the between and within components of the non-core dispersion response, quarters one to four, with one common scale concentrated out by weighted least squares (the total is their sum in data and model and is therefore excluded from the objective

Table J.9: Robustness of the inversion to the expenditure share

Calibration	ω_1	$\phi^*_{U,1}$ (food/gas)	$\phi^*_{U,2}$ (services)
Baseline	0.27	0.003	0
Equal sector shares	0.50	0.005	0
Salience-weighted perceived basket ($\omega_U = 0.57$)	0.27	0.003	0

Notes: Welfare-maximizing sectoral uncertainty responses under core targeting ($\theta_1 = 0$), from the second-order pruned exact-utility welfare object with the independent dispersion shocks active; each optimum is a sharp interior point on the reported grid. In every configuration the optimum leans on food-and-gas uncertainty ($\phi^*_{U,1} > 0$) and looks through services uncertainty ($\phi^*_{U,2} = 0$): the inversion is unchanged when the food-and-gas expenditure share is doubled from its empirical value to one half, and unchanged under the salience-weighted perceived basket, which reweights only the aggregate uncertainty object and therefore leaves the sectoral rule's optimum untouched. The result reflects the flexibility gap $\alpha_1 < \alpha_2$ and the much larger steady-state variability of food-and-gas inflation uncertainty ($\bar{h}_1 \gg \bar{h}_2$), not the expenditure weighting.

as an adding-up identity); the gas–food relative-price response scored on shape with its own concentrated scale; and the quarterly persistence of the measured uncertainty index, 0.88, with its delta-method standard error (0.07, from the monthly autocorrelation and $\rho_q = \rho_m^3$). Held fixed are the externally identified parameters: the sectoral rigidities (α_1, α_2), the policy rule, trend inflation, preferences, the survey underreaction rigidity, the between weight b_{g2} (a measured variance ratio), and the gap's unconditional volatility (an imposed identity that solves the innovation scales given ρ_d and the activity share).

Numerically, the general-equilibrium block is invariant to θ at the estimation configuration, so the model moments are computed from a single perturbation solution: the uncertainty responses through the exact linear filter mapping signal noise into posterior variance (fitted once on the reference solution, with near-perfect fit), the persistence from a long simulation of the same filter including the normalized between input, the between component by averaging the squared-gap perturbation over the model's ergodic gap distribution (simulated generalized impulse responses of squared states are noise-dominated at feasible replication counts), and the deferral–catch-up component in closed form from the policy-rate path. The objective evaluates in well under a second, the estimation converges to the same point from dispersed starting values, and profiling confirms the activity-share result. Three parameters sit at economically meaningful bounds and are reported as constrained estimates: the activity share of the gap at its floor of zero, the catch-up ratio $\nu_{\text{catch}}/\nu_{\text{defer}}$ at its cap of 0.90 (deferred adjustments release at most nine-tenths of the accumulated backlog), and ρ_d at its persistence floor of 0.20 (a minimum half-life for the vehicle-stock state); the volatility-factor sensitivity ψ_v is an interior estimate. The size of the services response is fitted through the salience loading κ_{serv}, identified by the services uncertainty moments; its timing is not parameterized, and the model's services trough arrives in the third quarter as in the data, at -1.9 against -2.3, with a trough ratio of 0.42 against 0.40.

Inference and weight robustness are assessed in four ways. First, a moving-block bootstrap over the BLS item panel (five hundred replications, fifteen months blocks, re-estimating the dispersion and gap projections jointly per draw) delivers the joint covariance of the BLS-based moments; the bootstrap standard errors exceed the independence-aggregated ones by factors of 1.4 to 1.8, confirming that adjacent-horizon projection

coefficients are strongly correlated, and the same factor is applied to the survey-based moments. Second, re-estimating with the bootstrap weights moves every parameter by a fraction of its standard error (κ_{disp} from 2.08 to 2.09, κ_{serv} from 0.83 to 0.84, ρ_q from 0.55 to 0.54) and leaves the activity-share result unchanged; the deferral–catch-up pair is weakly identified, with only the location of the dispersion response's sign change disciplined by the data, and we report it as such. Third, estimating the survey composites' projections directly at the quarterly frequency, with standard errors clustered by quarter, yields responses consistent with the aggregated monthly targets within their bands; the monthly projections remain the targets because the rotating panel supports them with far more within-respondent depth. Fourth, delta-method standard errors from the bootstrap moment covariance, holding the bounded parameters at their bounds and re-concentrating the fitted scales at each perturbation, are 2.0 for κ_{disp}, 1.7 for κ_{serv}, and 0.22 for ρ_q (0.93 and 19 for the deferral persistence and the catch-up intensity, and 793 for the volatility-factor sensitivity); the two dispersion loadings are estimated essentially uncorrelated, as they are identified by different sectors' moments. The bands are wide, consistent with the borderline significance of the underlying local projections, so we treat the point estimates as a device for quantifying the mechanism rather than as tightly measured structural constants; the robustness battery re-optimizes the rule across pass-through values spanning this range, and the positive optimum survives at every point. Within each evaluation the latent measurement constants $(\bar{h}_j, \bar{U}, \bar{\rho}_{12})$ are pinned by a one-shot fixed-point pass that iterates the square-root Kalman steady state against the Dynare `resol`-implied posterior moments until they are mutually consistent with the shock block and the pass-through κ; because the matched object is the relative-percent uncertainty response, it is invariant to the variance-units convention and to channel-toggle counterfactuals. Within each evaluation κ denotes κ_{disp}.

J.14 Robustness: Updated Targets, Agent-Specific Uncertainty, and Salience

This subsection records three robustness exercises on the policy-transmission calibration and the optimal-policy results, all solved with the exact nonlinear model at second order with pruning.

Robustness of the estimated transmission and of the leaning. The dispersion pass-through κ_{disp} is estimated on the household-survey uncertainty impulse responses of Section 5. The categories are expenditure-weighted into the two model sectors at the respondent level (food and gas into Sector 1; medical, education, and rent into Sector 2), so cross-category covariances are embedded in the composite; the projections are estimated on the composites and, because they are monthly while the model is quarterly, averaged within quarters. The estimated pass-through is $\kappa_{\text{disp}} = 2.08$, with the salience loading $\kappa_{\text{serv}} = 0.83$ (Appendix J.13). The prescriptions do not hinge on this value or on the rest of the policy block: Table J.10 re-optimizes the food-and-gas uncertainty response across the inflation coefficient, interest-rate smoothing, an output-gap response, the firms' information set, the pass-through itself, the sectoral rigidities, trend inflation, and the belief-noise settings (shutting off the memory, $\rho_q = 0$, and adding residual noise, $\sigma_h > 0$). The optimum is positive and the inversion unchanged in every configuration but one: at $\alpha_1 = 0.80$, where food and gas prices are made nearly as sticky as services, the optimum is exactly zero, because policy can then compress price-change dispersion

nowhere. That single exception is the pure mechanism test: the leaning comes from the flexibility gap, not from a generic response to uncertainty. Two configurations that eliminated the leaning at the pre-coupling calibration no longer do: with an output-gap response ($\phi_x = 0.125$) the optimum is unchanged with a larger gain, so the baseline choice $\phi_x = 0$ is a convention rather than a requirement, and without belief memory ($\rho_q = 0$) the optimum remains positive, because the salience channel transmits dispersion into both sectors' uncertainty without requiring accumulation.

Table J.10: Robustness of the leaning on food-and-gas uncertainty

Perturbation	$\phi^*_{U,1}$	CE gain (%)
Baseline ($\kappa_{\text{disp}} = 2.08$, estimated)	0.003	0.247
$\phi_\pi = 1.5$	0.003	—
$\phi_\pi = 3.0$	0.004	0.222
$\rho_i = 0$	0.003	—
$\rho_i = 0.7$	0.004	0.323
$\phi_x = 0.125$	0.003	0.401
Firms half informed	0.003	0.321
Firms fully informed	0.003	0.523
$\rho_q = 0$ (memory off)	0.003	0.285
$\sigma_{h,1} = 0.1$	0.003	0.246
$\kappa_{\text{disp}} = 1.0$	0.004	0.451
$\kappa_{\text{disp}} = 2.5$	0.004	0.184
$\alpha_1 = 0.60$	0.002	1.169
$\alpha_1 = 0.80$	0	0
$\alpha_2 = 0.85$	0.003	0.121
$\alpha_2 = 0.95$	0.003	—
Trend inflation 1%	0.002	0.247
Trend inflation 4%	0.020	—

Notes: Welfare-maximizing food-and-gas uncertainty response $\phi^*_{U,1}$ under core targeting ($\theta_1 = 0$, $\phi_{U,2} = 0$), re-optimized separately for each perturbation of the baseline, from the second-order pruned exact-utility welfare object with the independent dispersion shocks active. The CE gain is the consumption-equivalent gain of the optimal response over the uncertainty-blind rule ($\phi_{U,1} = 0$) in the same configuration. Dashes: for $\phi_\pi = 1.5$, $\rho_i = 0$, $\alpha_2 = 0.95$, and 4% trend inflation the configuration is close to the determinacy boundary and the consumption-equivalent conversion is unreliable, but the optimum remains positive and interior. At $\alpha_1 = 0.80$ the optimum is exactly zero: with food and gas nearly as sticky as services, policy cannot compress dispersion anywhere and the leaning shuts down, confirming that the flexibility gap is the mechanism; it is the battery's single exception. The optima are sharp interior points on the reported grid rather than plateaus. The optimal services response is zero at the baseline and in every basket configuration of Table J.9, and reacting to services uncertainty is strictly welfare-reducing at the optimum, because services uncertainty is dominated by components policy cannot compress.

Household versus firm uncertainty. The SCE measures household uncertainty, so we ask whether the results depend on firms perceiving the same uncertainty. We split the information share between households and firms, so that firms condition on the latent state with weight $\theta^{firm} \in \{0, \frac{1}{2}, 1\}$ rather than on the noisy posterior mean alone (the baseline sets $\theta^{firm} = 0$, shared beliefs). The decomposition in Appendix J.7 shows that at the estimated parameters the welfare cost of noisy information splits 97 percent firm-side and 3 percent household-side: firms price on the shared noisy beliefs, so household

uncertainty costs welfare mainly through their resets. Re-optimizing the food-and-gas uncertainty response as firms move from fully uncertain to fully informed keeps the optimum positive throughout ($\phi^*_{U,1} > 0$): the consumption-equivalent gain from leaning rises from 0.25 percent when both agents share the noisy beliefs to 0.32 and 0.52 percent as firms become half and fully informed, while the total cost of noisy information is 0.131 percent of consumption. The food-and-gas uncertainty prescription and the look-through-the-level result hold across all three cases, including when firms, which households own and manage, infer inflation far better than the average household. Because this lever changes firms' information about the latent state rather than the pricing-risk variance wedge, it isolates the role of firm information; the aggregate and sectoral prescriptions are robust to it.

Salience of food. Households overweight frequently purchased, salient categories when forming inflation beliefs, so one might worry that the results hinge on the weight the bank assigns to food and gas in an aggregate uncertainty index. We replace the true expenditure share ω_1 with a perceived weight ω_U in the basket-uncertainty object U_t, calibrating $\omega_U = 0.57$ so that the model's food-and-gas share of steady-state basket uncertainty matches the empirical food-and-gas contribution share from the decomposition of Section 4, about twice the budget share; the true ω_1 continues to govern the resource constraint, the price index, and welfare. The optimal rule is unchanged ($\phi^*_{U,1} = 0.003$, $\phi^*_{U,2} = 0$; Table J.9), and for a simple reason: the optimized rule is sectoral, responding to food-and-gas uncertainty directly, so no basket weight enters it. Salience reweighting only tilts an aggregate uncertainty index toward food and gas, which is the direction the optimal sectoral split already takes to its extreme (Figure 13).

Is the inversion a relabelling of the inflation response? Because food-and-gas uncertainty might co-move with food-and-gas inflation, a natural concern is that leaning on h_1 while looking through π_1 is a back-door response to π_1 rather than a genuine second-moment response. In the model the two are essentially orthogonal to begin with: the ergodic correlation between h_1 and π_1 is about -0.11, because the dispersion that drives the posterior variance accumulates through the mean-gap state rather than tracking current inflation. We nonetheless test the concern directly. We project π_1 out of the uncertainty object, forming $h_1^{\perp} = (h_1 - \bar{h}_1) - \beta\,(\pi_1 - \bar{\pi})$ with $\beta = \mathrm{cov}(h_1, \pi_1)/\mathrm{var}(\pi_1)$, so that $\mathrm{corr}(h_1^{\perp}, \pi_1) = 0$ by construction, and re-optimize the rule's uncertainty response on $h_1^{\perp}$ in place of h_1. The welfare gain from leaning on uncertainty is undiminished: reacting to the inflation-orthogonal component delivers essentially the same consumption-equivalent gain as reacting to raw h_1 (0.242 versus 0.247 percent, 98 percent surviving), at the same optimal response coefficient. The inversion is therefore a response to the component of food-and-gas uncertainty that moves independently of food-and-gas inflation, not a relabelled inflation response.

Volatility scale and the size of the welfare stakes. The welfare numbers are absolute consumption equivalents, and Table 6 shows the model's realized food-and-gas inflation volatility (3.0 percentage points) sits below the measured 4.8, so the stakes may be understated. We verify this directly with a homogeneous volatility rescaling: the latent-state volatilities are multiplied by a factor L and the signal-noise levels by L^2, and the belief steady state is recomputed through the filter's fixed point. The rescaling is exact in the belief block, with the steady-state posterior covariance scaling by L^2

and the log uncertainty levels shifting by $2 \ln L$, and it leaves the uncertainty impulse responses to the policy shock invariant to machine precision, so the estimated transmission parameters, the moment fit, and every relative object in the paper are unchanged by construction. At $L = 1.59$ the model matches the realized food-and-gas volatility level (4.8 percentage points; the services level rises to 0.44 against 0.79, since a common factor cannot move the sectoral ratio). The welfare stakes grow faster than the L^2 of the belief-error variance, because the price-change dispersion that drives signal noise fluctuates with inflation volatility around a mean pinned by trend inflation, so the uncertainty fluctuations themselves also scale: at the level-matched point the optimum is unchanged and the consumption-equivalent gain from leaning rises from 0.25 to 0.75 percent, about threefold. The welfare magnitudes reported in the paper are therefore conservative with respect to the volatility scale, while the prescriptions do not depend on it.

Linear versus generalized impulse responses. The uncertainty impulse responses used in the estimation are first-order responses, the natural counterpart to the linear local projections they are matched to. We do not report simulated generalized impulse responses of the uncertainty states: their sampling error at feasible replication counts exceeds the impact response itself (the ergodic standard deviation of log posterior variance is roughly forty times the impact response, so even thousands of replications leave the estimated correction dominated by noise), and estimates of the correction are accordingly unstable across parameterizations. The order robustness of the results does not rest on that object: the belief variances are carried as explicit states rather than recovered as perturbation residuals, so the welfare ranking is evaluated at second order directly (Appendix J.6), and the measured-dispersion components that are genuinely second-order are computed by averaging over the model's ergodic distribution in closed form rather than by simulation (Appendix J.4).

J.15 Determinacy and Pruning Diagnostics

Section 6.3 of the main text reports an interior optimum in the sectoral uncertainty responses. This subsection documents the determinacy checks behind that result. The results are also robust to the approximation order: re-solving the model at third order with pruning and re-evaluating the $\phi_{U,1}$ grid under common random numbers leaves the optimum at $\phi^*_{U,1} = 0.003$ and the welfare profile unchanged point by point, with the blind-rule cost moving from 25 to 23 basis points, so the second-order solution is adequate for the welfare ranking; the third-order terms it omits are the skewness wedges, whose analytical form Appendix J.9 derives and whose numerical contribution is small. We follow the pruning approach of Andreasen et al. (2018), who show that pruning yields a stationary, ergodic state-space at higher orders even when the unpruned non-linear solution diverges; spurious explosions are therefore distinct from Blanchard–Kahn (BK) indeterminacy at the linearized solution.

The diagnostic varies the two sectoral Taylor-rule loadings of the shipped rule (18) over $\phi_{U,1} \in \{0, 0.002, 0.005, 0.01, 0.02, 0.035, 0.06, 0.1\}$ and $\phi_{U,2} \in \{0, 0.005, 0.02\}$ at the baseline policy $(\rho_i, \phi_\pi, \phi_x) = (0.3, 2, 0)$ and the estimated pass-through $\kappa_{\text{disp}} = 2.08$, and records the BK status of the linearized solution at each grid point. Because the signal-noise states respond to realized year-on-year price-change dispersion rather than to the lagged policy stance, the saturating real-rate-gap transform present in the code plays no role in the belief block and is left inactive throughout the exercise.

The model is determinate at every grid point: the sectoral rule satisfies the BK conditions for a food-and-gas uncertainty response up to at least $\phi_{U,1} = 0.1$, more than thirty times the optimum $\phi^*_{U,1} = 0.003$, and for a services response up to at least $\phi_{U,2} = 0.02$, in every combination. Under positive trend inflation the determinacy region is more restrictive than in the zero-trend benchmark, which is why the baseline sets a zero output-gap response; at the conventional inflation response $\phi_\pi = 2$ with $\phi_x = 0$ the interior optimum of Section 6.3 lies deep inside the determinate set. The welfare deterioration at aggressive uncertainty responses in Figure 12 is therefore an interest-rate-volatility cost inside the determinate region, not a determinacy failure.

Table J.11: Determinacy and pruned ergodic moments across the uncertainty loading ϕ_U and perturbation order

	$\phi_{U,2} = 0$	$\phi_{U,2} = 0.005$	$\phi_{U,2} = 0.02$
$\phi_{U,1} \in \{0, 0.002, 0.005, 0.01, 0.02, 0.035, 0.06, 0.1\}$	determinate	determinate	determinate

Note: BK status of the linearized solution (Dynare `resol` info code = 0 at every grid point), evaluated for the sectoral rule (18) at the baseline policy $(\rho_i, \phi_\pi, \phi_x) = (0.3, 2, 0)$ and the dispersion pass-through $\kappa_{\text{disp}} = 2.08$. The grid extends to more than thirty times the optimal food-and-gas response and covers every combination of the two sectoral loadings. Under the trend-inflation dispersion mechanism the signal-noise states respond to realized price-change dispersion rather than to the policy stance, so the saturating real-rate-gap transform present in the code is inert and is not varied.