Using density forecasts to measure the stability of inflation expectations*

Ian Dew-Becker, Stefano Giglio, and Pooya Molavi

September 11, 2026

Abstract

The anchoring of inflation expectations is a major objective of modern monetary policy. Under very weak conditions, allowing for a wide range of deviations from the usual linear-Gaussian rational expectations structure, the local sensitivity of an agent's expectations to news is equal to their subjective uncertainty multiplied by the subjective precision of the observed signal. Instrumental variables estimates show that uncertainty is a significant determinant of expectational anchoring for survey participants, and the data is consistent with the model's implication of proportionality – a hypothetical agent with zero uncertainty would ignore all signals. At the aggregate level, average uncertainty predicts the sensitivity of average expectations to realized inflation surprises, and with a magnitude consistent with the micro IV-based estimates. The results imply that as of 2025, household inflation expectations were 2–3 times more sensitive to news than prior to 2020.

Background and contribution

One of the primary goals of modern monetary policy is ensuring that inflation expectations remain well anchored in the sense that they are not too sensitive to news and therefore do not fluctuate very much. The core difficulty in measuring the strength of anchoring is that we only observe expectations in the data conditional on the news that was actually realized. Without further assumptions, it is not possible to estimate the slope of the response of

*Dew-Becker: Federal Reserve Bank of Chicago; Giglio: Yale University; Molavi: Northwestern University. The views in this paper are those of the authors and do not represent those of the Federal Reserve Bank of Chicago. We appreciate discussions and comments from Jeff Campbell, Naveen Gondhi, and participants at the INSEAD finance symposium, SED, and NBER summer institute.

expectations to news because it requires observing counterfactuals – how people would have responded to alternative scenarios.[1]

This paper's approach starts from the observation that under weak conditions, the change in an agent's expectations in response to information is proportional to their current uncertainty multiplied by the precision of that information. It then asks the extent to which that formula describes the behavior of expectations in survey data both at the individual and aggregate levels.

The empirical analysis has two components: estimation and testing. Estimation corresponds to measuring the quantitative relationship between anchoring and reported uncertainty. By testing we mean evaluating the model's prediction that anchoring is *proportional* to uncertainty – that the intercept is zero, so that a hypothetical person with zero uncertainty has perfectly anchored expectations. We estimate and test the model in both micro and aggregate data.

Methods

The paper's approach is to model inflation expectations as coming from a filtering problem in which agents observe signals about future inflation, which is the latent state (a recent example of such a setup is Carvalho, Eusepi, Moench and Preston (2023)). If agents act under the assumption that signals are Gaussian (which can be motivated, for example, by assuming that they observe a large number of weakly informative signals with arbitrary distributions, or that they use normality as a simple rule of thumb), the Bayesian update step in the filter – the mapping from signals to expectations – has a tractable form. Dytso, Poor and Shitz (2022) show how to express a single Bayesian update as a power series, and Dew-Becker, Giglio and Molavi (2026) show how to implement the filter dynamically.

We use two key filtering results. First, again, the local sensitivity of expectations to signals is equal to uncertainty multiplied by signal precision. Second, the evolution of uncertainty itself depends on the realization of the signals multiplied by signal precision and the *third* moment, which helps formally characterize the sources of endogeneity in uncertainty.

The paper tests the model predictions using data from the New York Fed's *Survey of Consumer Expectations* (SCE). The SCE is particularly useful here because it has data available at the monthly frequency and asks respondents to provide probability distributions for future inflation.

In the cross-section of the SCE, the raw relationship between uncertainty and responsiveness

[1] See Ball and Mazumder (2011) for a discussion of anchoring in terms of either levels of expectations being stable over time or being insensitive to shocks. See also Kumar, Afrouzi, Coibion and Gorodnichenko (2015) and the discussion in Armantier et al. (2022)

to signals turns out to be if anything *negative*. That result is not surprising since uncertainty is endogenous to signal precision. Agents who consume a lot of information and therefore have precise signals will also tend to have the lowest uncertainty, making the reduced-form relationship between sensitivity and uncertainty ambiguous.

We therefore use the panel dimension of the data to run an instrumental variables analysis. Kim and Binder (2023) show uncertainty systematically declines with tenure in the survey and we give conditions under which it is a valid instrument. The decline in uncertainty is consistent with people on average paying more attention to inflation while they are in the survey than they did previously. It is natural to think that after being asked many questions about inflation, news about it would be more salient, and the questions in the survey themselves contain implicit information about inflation (e.g. its typical scale).

The treatment used here for identification is the survey itself. Instead of comparing treated to untreated people, though, the identification comes from *variation in the duration of the treatment*. People who have been in the survey longer have had increased attention for longer, making their uncertainty lower compared to people who have just entered the survey. The exogeneity assumption then is that attention is mean-independent of tenure – that is, while people pay more attention while in the survey than they did previously, we need the restriction that their attention does not decline over the course of the survey. The availability of this instrument is another feature of the SCE that distinguishes it from other surveys, such as the Survey of Professional Forecasters.

In the cross-section of the SCE, the reduced-form finding is that the survey respondents with low tenure have relatively high uncertainty and their expectations both covary more strongly with inflation surprises and are more volatile overall than those of respondents with high tenure, consistent with the model. Under the instrumental variables interpretation, then, the second stage says that high uncertainty is associated with larger responses to inflation surprises and greater volatility in beliefs: respondents with low tenure have expectations that are more weakly anchored than those with high tenure under both the level and shock concepts of Ball and Mazumder (2011).

The quantitative relationship between responsiveness to news and uncertainty – a slope – measures in the model agents' average signal precision. The model has an additional prediction, though, which is that an agent with zero uncertainty should have zero responsiveness – the intercept is zero. We show that the data does not reject that null.

We also examine time-series regressions using average uncertainty and inflation expectations. The magnitude of the aggregate time-series relationship between average uncertainty and the sensitivity of average expectations to inflation is consistent with the micro IV-based

estimates, and we again get the result that the estimated intercept is zero.

Related work

Armantier et al. (2016) provide closely related evidence from an information provision experiment that agents are approximately Bayesian in the sense that their responsiveness to information increases in prior uncertainty.[2] Related work includes Weber et al. (2025), Coibion and Gorodnichenko (2015), Cavallo et al. (2017), and Gemmi and Mihet (2026). Note that those papers work in rational linear-Gaussian setups, which the theoretical structure here generalizes significantly.

Another complementary approach is Armantier et al. (2022), which directly asks people about how their expectations would change in response to certain specific events, such as particular realizations of inflation or unemployment.[3]

Lamla and Vinogradov (2019) and Binder, Campbell and Ryngaert (2024) also study responses of inflation expectations in the SCE, focusing on within-month dynamics (using the fact that the SCE reports the exact date on which each response was recorded).

Carvalho et al. (2023) is also very closely related. It posits a New Keynesian structure for inflation dynamics and then assumes that agents form expectations through a rule of thumb filtering rule. The inflation process in that paper is one particular case of the general class of processes allowed in this paper's analysis. The primary contrast is that this paper assumes that agents are Bayesian or quasi-Bayesian, as opposed to boundedly rational.

Given this paper's focus on the conditional distribution of future inflation, it is naturally also related to work that studies option-implied distributions, including Kitsul and Wright (2013), Fleckenstein, Longstaff and Lustig (2017), Mertens and Williams (2021), and Hilscher, Raviv and Reis (2025). The basic implication of this paper's analysis is that option-implied uncertainty would also be a measure of expectational anchoring.

More generally, Reis (2025) reviews recent evidence on inflation expectations and their link to post-covid inflation. See also D'Acunto and Weber (2024) and Coibion and Gorodnichenko (2025).

Outline

Section 1 develops the paper's core economic structure and the link between posterior beliefs and the response function for expectations. Section 2 examines how the theory can be tested empirically and section 3 provides estimates. Finally, section 5 studies the model's implications for forecasting expectational anchoring going forward and section 6 concludes.

[2] See also Weber et al. (2025).

[3] Ameriks et al. (2011) introduce the idea of strategic surveys as a solution to the sort of identification problem that this paper faces. When counterfactuals are not observable, the approach is to simply ask about them. Armantier et al. (2015) and Roth and Wohlfart (2020) among others provide evidence that people act on their beliefs.

1 Theoretical framework

This section characterizes the drivers of expectations under relatively weak assumptions about priors and rationality. π_t^* denotes the object agents are learning about (either trend inflation or actual realized future inflation) and y^t their history of signals. Their date-t mean and variance are then

$$m_t \equiv E\left[\pi_t^* \mid y^t\right] \tag{1}$$

$$u_t \equiv \text{var}\left(\pi_t^* \mid y^t\right) \tag{2}$$

1.1 Basic updating result

1.1.1 Economic environment

Assumption 1 *π_t^* follows an arbitrary dynamic process. It is not necessarily linear, Gaussian, or homoskedastic.*

Empirically, US inflation historically is positively skewed (the skewness coefficient for core PCE inflation is 1.24) and its volatility has varied significantly over time (e.g. Engle (1982) and Stock and Watson (2007)), so it is preferable if possible to not impose normality on its dynamics.

While we allow general priors, signals take the usual Gaussian form.[4]

Assumption 2 *Agents receive a signal y_t about π_t^* which they assume is distributed as*

$$y_t \sim N\left(\pi_t^*, \sigma_t^2\right) \tag{3}$$

σ_t contains no information about future inflation beyond what is contained in y_t itself.

σ_t^{-2} can be interpreted as measuring the attention agents pay to inflation. In practice the signal is agent-specific, and below we add an i subscript to indicate different agents. If individual agents receive multiple normally distributed signals, they will have a sufficient statistic of the form (3).

[4]One way to motivate the assumption of normality is to assume that agents receive many independent signals centered on π_t^*, each with very low precision, that have arbitrary distributions. Then the martingale central limit theorem implies that the sum of the signals (i.e. y_t) is asymptotically – as the precision goes to zero and the number of signals goes to infinity at the same rate – normally distributed and the sufficient statistic for the Bayesian update (Hall and Heyde (2014)). Alternatively, agents might truly receive normally distributed signals, or, again, they might assume normality as a simple rule of thumb.

There are other versions of the main result we use under alternative distributional assumptions for the signal. Normality is the usual baseline, but in other settings a different distribution (e.g. Poisson) could be natural.

1.1.2 Bayes' formula but not necessarily rational expectations

Assumption 3 *On all dates agents have subjective probability distributions over potential values of π_t^*. The distributions are updated using Bayes' formula based on their beliefs about the dynamic process driving π_t^* and the signal distribution (assumption 2).*

Assumption 3 says that just that agents use Bayes' *formula*. It does not impose that their specification for either the signals (assumption 2) or for the dynamics of π_t^* (assumption 1) is correct. Specifically, we can have the following deviations from rationality:

1. **Signal distribution:** Agents may use an incorrect value for signal precision (σ_t^{-2}) in calculating their update. They might place too much or too little weight on their signals (e.g. as in diagnostic expectations) or ignore some sources of information. The signals might also not actually be Gaussian.

2. **Prior distribution:** Agents' prior need not be correct in any way. It might be wrong if agents have the wrong dynamic model for inflation, they might have some strange prior based on superstition, or they might just for some reason have started with beliefs far from the truth (e.g. Farmer, Nakamura and Steinsson (2024)). The prior need not even be absolutely continuous with respect to any sort of true distribution, and agents need not have common priors (e.g. as in Malmendier and Nagel (2016)).

What the analysis imposes is simply that agents act as though their signals have normally distributed errors and that they update using Bayes' formula. Many behavioral models are therefore consistent with the paper's setup, including Daniel et al. (1998) and DeMarzo et al. (2003) (see Molavi (2025)). It is also straightforward to show that the paper's results hold under diagnostic expectations (as in Bordalo et al. (2019)), with the behavioral bias appearing as an effective rescaling of the signal precision.

1.1.3 Updating expectations

The standard updating formula for mean expectations that holds in the fully Gaussian case holds locally here.

Proposition 1 *Under assumptions 1–3, in the limit when time periods are short or signals are uninformative, expectations are updated according to*

$$\Delta m_t \approx u_t \sigma_t^{-2} \left(y_t - m_{t|t-1} \right) + E_{t-1} \left[\Delta m_t \right] \tag{4}$$

where $m_{t|t-1} \equiv m_{t-1} + E_{t-1} \left[\Delta m_t \right]$.

Proposition 1 obtains as a discrete-time approximation from the results of Dew-Becker, Giglio and Molavi (2026).[5]

(4) is the main equation that the paper aims to test. The true update is nonlinear, but to first order the slope is equal to agents' uncertainty multiplied by their belief about signal precision.[6]

1.1.4 Dynamics of uncertainty

The gain in equation (4) is $u_t\sigma_t^{-2}$. It is fairly common to measure uncertainty empirically, while signal precision is generally treated as unobservable. In comparing agents' sensitivity to news to their uncertainty, the obvious concern is that signal precision might be correlated with uncertainty.

Proposition 2 *Under assumptions 1–3, in the limit when time periods are short or signals are uninformative, uncertainty updates according to*

$$\Delta u_t \approx \text{var}_{t-1}\left(\Delta\pi_t^*\right) + 2\,\text{cov}_{t-1}\left(\Delta\pi_t^*, \pi_t^*\right) - \sigma_t^{-2}u_{t-1}^2 + \sigma_t^{-2}\mu_{3,t}\left(y_t - m_{t|t-1}\right) \tag{5}$$

where $\mu_{3,t}$ denotes the third central moment of the posterior distribution of π_t^*.

This again obtains as a discretization of a result from Dew-Becker, Giglio and Molavi (2026).[7]

Proposition 2 generalizes the update equation for uncertainty that holds in the normal prior/normal signal case, with only the final term in (5) being different. The first two terms represent the increase in uncertainty from shocks to fundamentals. The third term is the decline coming from information accumulation and the fourth is the effect of signals on uncertainty.

There are two important implications of proposition 2 for the analysis here. First, there will in general be a negative cross-sectional relationship between u_t and σ_t^{-2} (assuming σ_t^{-2} is persistent), which will flatten or even reverse the relationship between sensitivity to news and uncertainty.

[5] Dytso, Poor and Shitz (2022) show how to write the posterior expectation as a power series in terms of the signal. The coefficients in the power series are the posterior cumulants. The approximation error in (5) comes from those terms. Specifically, the next order term is $\frac{1}{2}u_t^2\sigma_t^{-4}\left(y_t - m_{t|t-1}\right)^2$. Proposition 1 is the discretization of proposition 2 in Dew-Becker, Giglio and Molavi (2026) with the parameter a equal to zero.

[6] In principle the approximation can be written in terms of either u_t or u_{t-1} – the difference is second-order (and in continuous time they are equal). We write it in terms of u_t because that maps directly onto the analogous discrete-time result (see Dytso, Poor and Shitz (2022)). Formally, $\partial m_t/\partial y_t = u_t\sigma_t^{-2}$.

[7] As above, the error in this approximation involves the higher-order terms in a power series. The next order term is $\frac{1}{2}\kappa_{4,t}\sigma_t^{-4}\left(y_t - m_{t|t-1}\right)^2$, where $\kappa_{4,t}$ is the fourth posterior cumulant.

Second, u is endogenous to the actual realization of the signals with a coefficient proportional to the third moment (again, the posterior is used because, formally, $\partial u_t/\partial y_t = \sigma_t^{-2}\mu_{3,t}$). When beliefs are asymmetrical, signals affect uncertainty by determining whether agents think the longer or shorter tail of the distribution is more likely to be realized.

1.2 Observable predictions

This section derives the equations that the paper estimates empirically.

1.2.1 Squared changes in expectations

Squaring equation (4), taking expectations, we have

Corollary 1 *For an individual agent*

$$E_{t-1}\left[(\Delta m_t)^2\right] \approx \left(E_{t-1}\left[u_t\right]\right)^2 \sigma_t^{-2} \tag{6}$$

ignoring higher order terms in σ^{-2}.

Corollary 1 says that squared changes in beliefs are on average equal to squared (expected) uncertainty multiplied by signal precision.[8] In principle a regression of squared changes in expectations on squared uncertainty should potentially yield an estimate of average signal precision. Corollary 1 also implies that the coefficient in a regression of volatility in expectations on squared uncertainty should have an intercept of zero, which is a prediction that yields a test of the model.

1.2.2 Responses of expectations to realized inflation

While we do not know all the information agents use in forecasting inflation, it is natural to think that realized inflation is one component. To rationalize that and also impose some structure to help in estimation, a natural special case of the model is the trend-plus-noise specification for inflation.

Definition 1 ***Trend plus noise model.*** π_t^* *represents trend inflation. Realized inflation,* π_t*, follows*

$$\pi_t = \pi_t^* + \varepsilon_{\pi,t}, \text{ with } \varepsilon_{\pi,t} \sim N\left(0, \sigma_{\pi,t}^2\right) \tag{7}$$

[8] Here we have assumed that σ_t^{-2} is the true precision of the errors. If agents believe it to be σ_t^{-2} but the true precision is actually $\bar{\sigma}_t^{-2}$, then the equation becomes $E_{t-1}\left[(\Delta m_t)^2\right] \approx \left(E_{t-1}\left[u_t\right]\right)^2 \sigma_t^{-4}\bar{\sigma}_t^2$. Since we take precision as unobservable, this makes little difference to the analysis.

Agents observe a noisy signal about realized inflation, $s_{i,t}$ (consistent with limited attention or with people just observing a subset of prices):

$$s_{i,t} = \pi_t + \varepsilon_{s,i,t}, \text{ with } \varepsilon_{s,i,t} \sim N\left(0, \sigma^2_{s,i,t}\right) \tag{8}$$

Agents also observe a second signal, $x_{i,t}$, that is independent of $s_{i,t}$ conditional on π^*_t and has an arbitrary distribution. $x_{i,t}$ represents all the other information agent i gets about trend inflation (e.g. from policy announcements) and acts as a residual.

The trend plus noise model does not require either that inflation is perfectly observed by agents or that it is their only source of information. It just rationalizes why realized inflation might be one piece of relevant information.

In the trend plus noise model, proposition 1 can be adjusted to get the following result.

Corollary 2 *Define $\hat{\pi}_{i,t} \equiv \pi_t - m_{i,t-1}$ and $\sigma^{-2}_{i,t} \equiv \left(\sigma^2_{s,i,t} + \sigma^2_{\pi,t}\right)^{-1}$. Then in the trend plus noise model*

$$\begin{aligned} \Delta m_{i,t} &= u_{i,t}\sigma^{-2}_{i,t}\left(\pi_t - m_{i,t-1}\right) + \varepsilon_{i,t} && (9) \\ \text{where } \varepsilon_{i,t} &\equiv u_{i,t}\sigma^{-2}_{i,t}\left(s_{i,t} - \pi_t\right) + f_{i,t-1}\left(x_{i,t}\right) + \text{higher order terms} && (10) \end{aligned}$$

$f_{i,t-1}(\cdot)$ is a function representing how expectations are updated given the unobservable signal $x_{i,t}$.[9] *Up to terms of order $\sigma^{-2}_{i,t}$ or smaller, $u_{i,t}\sigma^{-2}_{i,t}\left(\pi_t - m_{i,t-1}\right)$ and $\varepsilon_{i,t}$ are both mean zero under the agent's probability measure.*

The trend plus noise model motivates regressions of changes in expectations on uncertainty multiplied by the inflation surprise. There is again a zero restriction, which here is that the sensitivity to inflation surprises should be zero for a hypothetical agent with zero uncertainty.

2 Empirical identification

This section discusses what assumptions are needed in order to estimate the regressions suggested by the theoretical analysis. Throughout the discussion, the assumption is that uncertainty is observable while signal precision is not.

[9] Specifically, $f_{i,t-1}\left(x_{i,t}\right) = E\left[\pi^*_t \mid y^{t-1}, x_{i,t}, \hat{s}_t = 0\right] - E\left[\pi^*_{t-1} \mid y^{t-1}\right]$.

2.1 Tenure as an instrument

The endogeneity of uncertainty to both $\sigma_{i,t}^{-2}$ and $\varepsilon_{i,t}$ implies that an instrument is necessary for identification. Motivated by the finding in Kim and Binder (2023) that uncertainty declines on average with tenure in the survey, we use tenure as the instrument.

Kim and Binder (2023) discuss how uncertainty may decline with tenure if people pay more attention to inflation news when they are in the survey than they did previously or if they even learn from the survey itself. In terms of the model, that corresponds to signal precision – $\sigma_{i,t}^{-2}$ – increasing on entry into the survey (more precisely, immediately after the first time they respond to the survey questionnaire), which then leads $u_{i,t}$ to decline over time via equation (5).

The treatment here is therefore the survey itself. Importantly, though, identification does not come from comparing treated to untreated agents – that is impossible since the method of observation is itself the treatment. Instead, it comes from variation in the *duration* of the treatment.

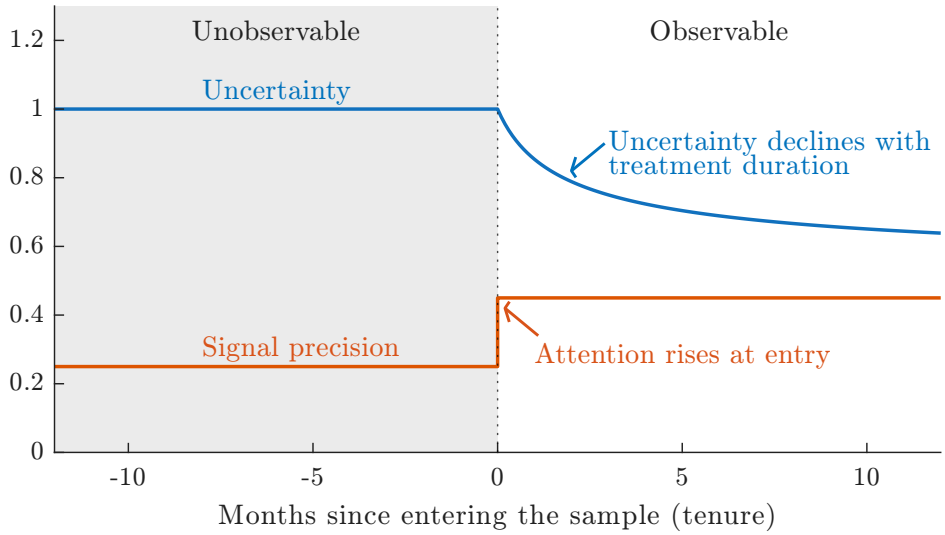
Figure 1 depicts the basic idea behind the identification. The x-axis represents the number of months since a respondent has entered the survey, with points to the left of zero corresponding to dates prior to the survey. The assumption that attention rises on entry to the survey appears here as a level shift in signal precision. That then causes their uncertainty to drift down through equation (5). The identification is not from comparing people pre- and post-entry to the survey, since we do not observe people prior to the survey. Instead, the identification in the regression comes from the the decline in uncertainty with tenure coming from higher attention. The basic prediction then is that the sensitivity expectations should fall with tenure.

In the figure, attention is depicted as rising immediately on entry to the survey. Whether the rise in attention happens while the person answers the first survey or in the time immediately after (when economic news is more salient) is irrelevant in the end because the usable observations in the regression are not until the second month of tenure, since the left-hand side involves the change in expectations (formally, see assumption 5).

2.2 Identifying assumptions

Because the endogeneity here comes in an interaction ($u_t\sigma_t^{-2}$), two-stage least squares in general need not yield a consistent estimator of the mean of $\sigma_{i,t}^{-2}$. While we report results for 2SLS, the primary results are for a control function (CF) estimator, which Wooldridge (2015) and Masten and Torgovitsky (2016) argue is preferable in settings such as ours with

Figure 1: Intuition behind the identifying assumption



Note: The figure shows signal precision and uncertainty for a hypothetical agent who satisfies the identifying assumption. They enter the survey at point 0 on the x-axis. The blue line represents their uncertainty, $u_{i,t}$, and the orange line the precision of their signals ($\sigma_{s,i,t}^{-2}$) – how much attention they pay to inflation. In the period that they are in the survey, the assumption is that attention is unrelated to tenure (though need not be constant), leading uncertainty to progressively decline.

heterogeneous slopes and in which the endogenous variable enters as an interaction.

The CF method still requires exogeneity and relevance assumptions, but stated in a slightly different way from in 2SLS.

Denote tenure by $z_{i,t} \in \{1, ..., 12\}$.

Assumption 4 *[Relevance] For $z_{i,t} \geq 2$ there is a first-stage representation for uncertainty, $u_{i,t}$,*

$$u_{i,t} = a_u + b_u z_{i,t} + \eta_{i,t}, \tag{11}$$

where tenure, $z_{i,t}$, is independent of the residual $\eta_{i,t}$, $b_u \neq 0$, and

$$\text{var}\left(z_{i,t} \mid \eta_{i,t}, \hat{\pi}_{i,t}\right) > 0 \tag{12}$$

The first stage need not necessarily be linear, but we take that as our benchmark as in 2SLS.[10]

[10]See Wooldridge (2015) and Masten and Torgovitsky (2016) for discussions of the first stage assumption.

Assumption 5 *[Exogeneity] Tenure, $z_{i,t}$, is exogenous and mean-independent of $\sigma_{i,t}^{-2}$ in the sense that*

$$\mathbb{E}\left[\sigma_{i,t}^{-2} \mid z_{i,t}, \eta_{i,t}, \hat{\pi}_{i,t}\right] = \mathbb{E}\left[\sigma_{i,t}^{-2} \mid \eta_{i,t}, \hat{\pi}_{i,t}\right] \quad (13)$$

$$\mathbb{E}\left[\varepsilon_{i,t} \mid z_{i,t}, \eta_{i,t}, \hat{\pi}_{i,t}\right] = \mathbb{E}\left[\varepsilon_{i,t} \mid \eta_{i,t}, \hat{\pi}_{i,t}\right] \quad (14)$$

for $z_{i,t} \geq 2$, where $\mathbb{E}$ denotes the population expectation operator.

In terms of the discussion from the previous section, attention rising on average on entry into the survey is what creates relevance because it creates a trend decline in uncertainty with tenure. The exogeneity assumption requires that attention is on average unrelated to tenure during the survey. If attention rises with tenure – e.g. as people see more iterations of the survey, inflation becomes more salient – that would counteract the effect of uncertainty and bias our estimates towards zero, so that the empirical results would represent a lower bound. On the other hand, if attention falls with tenure – e.g. the novelty of the survey wears off and people go back to their pre-survey level of attention – the estimates will be biased upwards, overstating the dependence of updating on uncertainty.

2.2.1 Is attention related to tenure?

One potential way to validate the estimates is to look at consistency between the micro estimates here and the time-series estimates in section 5, which use a fundamentally different identifying assumption that does not rely on tenure. We will show that in fact the micro and macro estimates have quantitatively similar implications.[11] If the two methods yield similar results, that suggests that even if attention is related to tenure, the bias may not be that strong in practice.

An alternative approach is to use the dynamics for uncertainty in equation (5) to back out an estimate of average signal precision relative to the pre-survey in each month of tenure based on the rate of decline in uncertainty. Appendix A.1 explains the calculation, and plots the result in figure A.4. Across four estimates (corresponding to different choices of which forecast from the survey to use and the forecast horizon), signal precision is approximately constant across months of tenure in two cases, increasing in one case, and decreasing in one case, suggesting that the assumption that it is constant is a reasonable description of the data.

[11] Within the survey itself there are no direct measures of attention, such as time to complete it. The closest is a question asking about enjoyment of the survey, which, for people who stay for all 12 months, does not display any relationship with tenure. We also show below that the main results are highly similar conditioning on the subsample that stays through all 12 months of the survey.

2.2.2 The second stage for the control function estimator

The control function estimator uses the residual from the first stage, $\eta_{i,t}$ in equation (11), as a control in the second stage. The basic second-stage regression for the trend-plus-noise model is

$$\Delta m_{i,t} = u_{i,t}\hat{\pi}_{i,t}\mathbb{E}\left[\sigma_{i,t}^{-2} \mid \eta_{i,t}, \hat{\pi}_{i,t}\right] + \mathbb{E}\left[\varepsilon_{i,t} \mid \eta_{i,t}, \hat{\pi}_{i,t}\right] + \nu_{i,t} \tag{15}$$

The two expectations play the role of the slope and intercept. Similarly, for the squared change in beliefs, the second stage is

$$(\Delta m_{i,t})^2 = u_{i,t}^2\mathbb{E}\left[\sigma_{i,t}^{-2} \mid \eta_{i,t}\right] + \mathbb{E}\left[\lambda_{i,t} \mid \eta_{i,t}\right] + \mu_{i,t} \tag{16}$$

The equations fit into the partially nonparametric setup analyzed by Newey and Stouli (2021), who show that the conditional expectation functions are identified (their theorem 4).

The fundamental parameter to be estimated is $\mathbb{E}\left[\sigma_{i,t}^{-2}\right]$. That can be estimated by averaging the conditional expectation across the values of $\eta_{i,t}, \hat{\pi}_{i,t}$. We estimate the conditional expectation functions by series least squares – i.e. approximating them with a set of basis functions (starting from a simple linear specification) – as suggested by Newey and Stouli (2021) (and see Newey (1997)).

2.3 Two-stage least squares

In general 2SLS does not yield consistent estimates of $\mathbb{E}\left[\sigma_{i,t}^{-2}\right]$.

Proposition 3 *In addition to the assumptions above, suppose $Z_{i,t} \equiv z_{i,t}\hat{\pi}_{i,t}$ is a valid instrument for $u_{i,t}\hat{\pi}_{i,t}$, in that $\mathbb{E}\left[Z_{i,t}\varepsilon_{i,t}\right] = 0$. Then the population second-stage estimate from a 2SLS regression of $\Delta m_{i,t}$ on $u_{i,t}\hat{\pi}_{i,t}$ and a constant with instrument $Z_{i,t}$ is*

$$\beta^{2SLS} = \frac{\mathbb{E}\left[\sigma_{i,t}^{-2}u_{i,t}\hat{\pi}_{i,t}\hat{Z}_{i,t}\right]}{\mathbb{E}\left[u_{i,t}\hat{\pi}_{i,t}\hat{Z}_{i,t}\right]} \tag{17}$$

where $\hat{Z}_{i,t} = Z_{i,t} - \mathbb{E}\left[Z_{i,t}\right]$.

As usual, the 2SLS estimate is a weighted average of the true treatment effect, in this case $\sigma_{i,t}^{-2}$. But the weights need not be positive here, meaning that the estimate β^{2SLS} need not even lie within the range of the $\sigma_{i,t}^{-2}$. A similar result holds for the regression with squared changes in expectations. That said, it is entirely possible for this average to happen to be near the population average, and we report conventional 2SLS results as robustness.

3 Data

We study the Federal Reserve Bank of New York's *Survey of Consumer Expectations* (SCE). The SCE is well suited to testing this paper's hypotheses both because it is a relatively large panel – with respondents remaining in the survey for up to 12 months – and because it is a representative sample of consumers, meant to capture inflation expectations of typical agents in the economy. The SCE is also run monthly, rather than quarterly, giving relatively fine time-series detail, which is useful since the most recent episode of high inflation lasted only about two years.

In the sample the monthly cohorts average 158 respondents. Respondents remain in the survey for 12 months, and after attrition there are on average 1,262 total observations per month.

The SCE asks respondents about inflation expectations at horizons of 1–12, 25–36, and 49–60 months. First, it asks for point estimates: "What do you expect the rate of inflation/deflation to be over the next 12 months? Please give your best guess." Next, it asks for probabilities that inflation falls into different bins: "what would you say is the percent chance that, over the next 12 months..." followed by a list of bins for inflation: 0–2%, 2–4%, 4–8%, 8–12%, and >12% along with the same bins for deflation. Equivalent questions are asked for annual inflation starting 24 and 48 months in the future. We use both the point estimates and the bins in the analysis.

For uncertainty we use the estimates of each respondent's variance that is constructed by the administrators of the SCE, who do so by fitting a beta distribution to the bin probabilities.[12]

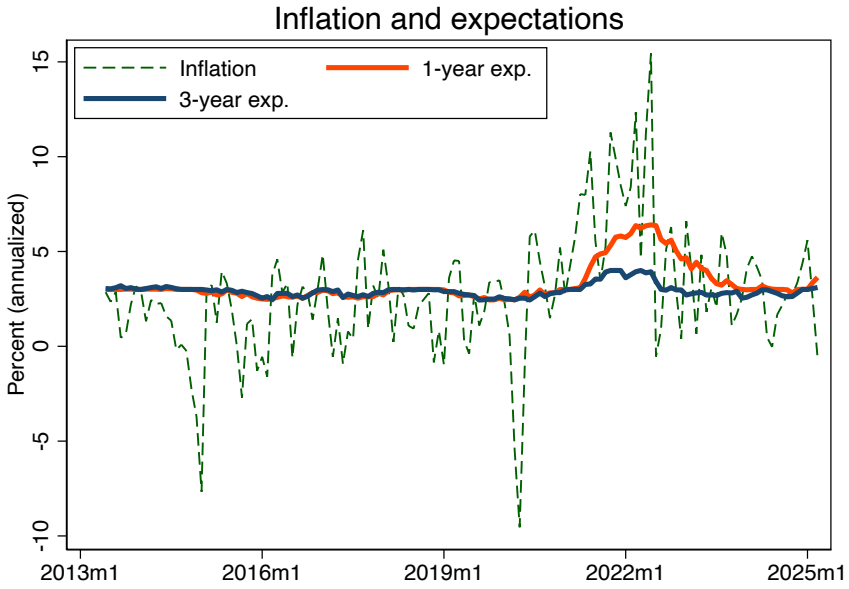
The paper's sample period is from the beginning of the SCE in June, 2013, through July, 2025. For context, figure 2 plots the time series of headline CPI inflation over that period along with 1- and 3-year inflation expectations (panel (a)) and with 1- and 3-year uncertainty (panel (b)), obtained from the SCE. Because the 5-year expectations were only included in the later waves of the survey, we do not include them in the analysis.

3.1 Mapping from the data to the model

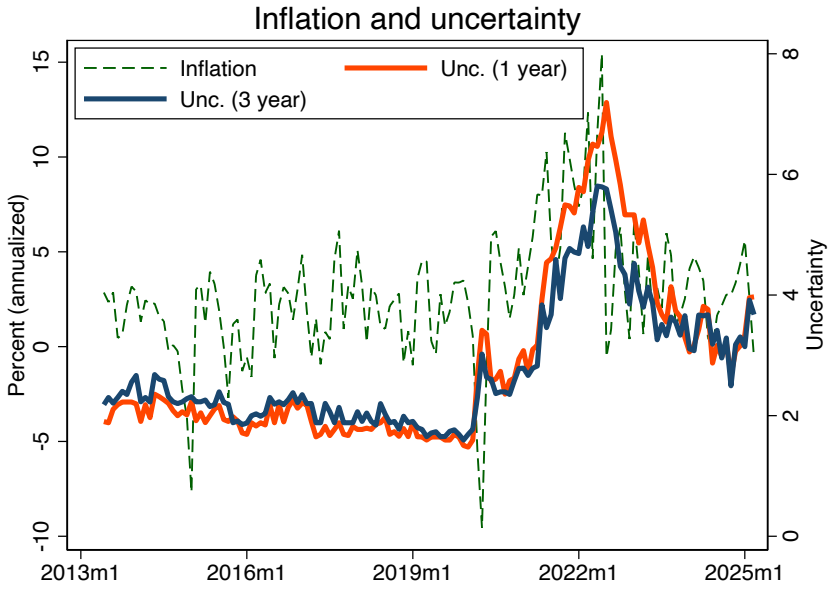
In the case of squared changes in expectations, π_t^* in the model corresponds to future realized inflation. The reported probability distributions for future inflation then correspond to

[12]They do not calculate a variance for responses that put positive probability on a set of non-contiguous bins, and we therefore drop those observations from the analysis. We additionally drop observations where the variance is greater than 200 squared percentage points and those with bin means outside the 5th and 95th percentiles or where the absolute change in expectations is greater than 8 percentage points.

Figure 2: Inflation, expectations, and uncertainty



(a) Inflation and expectations



(b) Inflation and uncertainty

Note: Panel (a) panel plots CPI inflation (annualized) together with the cross-sectional median 1-year and 3-year inflation expectations from the SCE. Panel (b) plots CPI inflation (left axis, annualized) together with the cross-sectional median uncertainty measures, at the 1-year and 3-year horizon on the right axis.

probability distributions for π_t^*.

For the trend plus noise model, where we study how expectations respond to realized inflation, π_t^* represents trend inflation. The survey, however, does not ask about the trend; it asks about the realization. For expectations, that difference is not a problem, since the noise in the model is mean-zero. Agents' subjective distribution over inflation, however, is not equal to their distribution for π_t^*. Instead, given equation (7), we have, at any horizon,

$$\text{var}_t\left(\pi_{t+j}\right) = \underbrace{\text{var}_t\left(\pi_t^*\right)}_{=u_t} + j\,\text{var}\left(\Delta\pi_t^*\right) + \sigma_\pi^2 \tag{18}$$

where here for simplicity we are assuming that the innovations to the trend and the noise component of inflation both have constant variance and that the trend is a random walk. The left-hand side corresponds to the variance actually reported in the survey, while the first term on the right is the relevant concept of uncertainty from the model. Under the assumption that the last two terms in (18) are constant over time, agents' reported subjective uncertainty over future inflation is equal to their uncertainty about trend inflation plus a constant.

Since all of the terms in (18) are nonnegative, reported uncertainty is zero only if u_t is also equal to zero, so a person who reports zero uncertainty in the trend plus noise model should have zero sensitivity to realized inflation.

3.2 Calculating inflation surprises

The benchmark results calculate the inflation surprise in each month for each person as realized inflation minus that person's lagged one-year expectation. Realized inflation can be measured as PCE or CPI, headline or core, while the lagged expectation can be either the point forecast or the distributional mean, giving us eight possible permutations. For the main text, the analysis uses CPI headline inflation. $\Delta m_{i,t}$ on the left-hand side of the regressions is measured based on the mean from respondents' probability distributions, and we report results using both the point forecast and bin mean for the expectations on the right-hand side. The remaining permutations are reported in the appendix.

The CPI and PCE inflation values for a given month are not available by the time of the survey. We use them to proxy for the information agents actually observe, including in particular the prices they pay for goods and services and examine robustness to that choice in section 4.5.

Since the change in expectations, $\Delta m_{i,t}$, is on the left-hand side of the regression, any time the same measure of expectations is on the right- and left-hand sides of the regression, the analysis uses the second lag of expectations in calculating the surprise to avoid contamination

from measurement error in the expectations. That is, for those cases $\hat{\pi}_{i,t} \equiv \pi_t - m_{i,t-2}$. Otherwise, $\hat{\pi}_{i,t} \equiv \pi_t - m_{i,t-1}$.

4 Results

4.1 OLS estimates

We begin with OLS results, i.e. without using the tenure instrument. The top two panels of table 1 reports results from regressions of $\Delta m_{i,t}$ on $\hat{\pi}_{i,t}$ and $\hat{\pi}_{i,t}u_{i,t}$ for the two inflation surprise measures (relative to the point forecast and distributional mean). Panel (a) reports results for 1-year expectations and panel (b) 3-year expectations. Columns (1) and (2) use contemporaneous uncertainty and columns (3) and (4) use lagged uncertainty. As one would expect, positive inflation surprises are associated with increases in future expectations. However, the interaction of the surprise with uncertainty is in most cases estimated to be *negative*, the opposite sign to what Bayesian updating predicts if the cross-sectional variation in uncertainty is exogenous. That is consistent with uncertainty being related to attention or to realized inflation.

In the regression, there are two sources of endogeneity: that uncertainty is affected by the same signals as expectations, and that uncertainty is correlated with signal precision. A way to distinguish them is to lag uncertainty in the regression, which should avoid the problem of uncertainty responding to the signals that drive $m_{i,t}$. Columns (3)–(4) do that and show that the coefficient on the interaction does rise, becoming either less negative or slightly positive.

The fact that the coefficients increase implies that signals that increase expectations also increase uncertainty, which the model attributes to positively skewed beliefs about inflation. At the same time, the fact that two of the coefficients remain negative implies that there is also still a remaining bias from uncertainty being correlated with signal precision in the cross-section.

Panels (c) and (d) report results for regressions of $(\Delta m_{i,t})^2$ on contemporaneous and lagged squared uncertainty. In this case, the interaction coefficients are positive, but quantitatively small. In particular, if they estimate $E\left[\sigma_{i,t}^{-2}\right]$, the coefficients of 0.0002–0.0008 imply that agents receive signals about future inflation each month whose errors have a standard deviation of 35–70 percentage points. As in the top two panels, the coefficients when using lagged uncertainty are larger than those using contemporaneous uncertainty.

Table 1: OLS reduced-form regressions of $\Delta m_{i,t}$ on inflation surprise and surprise × uncertainty (panels a–b) and of $(\Delta m_{i,t})^2$ on squared uncertainty (panels c–d)

(a) 1 year inflation

	Dep. var.: $\Delta m_{i,t}$			
	Contemp. uncertainty		Lagged uncertainty	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr.	−0.00124***	−0.00075***	−0.00062***	0.00037***
	[−0.00153, −0.00095]	[−0.00102, −0.00048]	[−0.00091, −0.00033]	[0.00010, 0.00064]
Infl. Surpr.	0.09405***	0.03716***	0.08756***	0.02471***
	[0.07676, 0.11134]	[0.02665, 0.04767]	[0.06992, 0.10520]	[0.01451, 0.03491]
Type of infl. surpr.	Point	Mean	Point	Mean
R^2	0.029	0.004	0.027	0.004
Obs.	103,888	93,460	103,888	93,460

(b) 3 year inflation

	Dep. var.: $\Delta m_{i,t}$			
	Contemp. uncertainty		Lagged uncertainty	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr.	−0.00076***	−0.00056***	−0.00050***	0.00017
	[−0.00100, −0.00052]	[−0.00080, −0.00032]	[−0.00073, −0.00027]	[−0.00008, 0.00042]
Infl. Surpr.	0.05313***	0.02524***	0.05045***	0.01715***
	[0.04130, 0.06496]	[0.01738, 0.03310]	[0.03865, 0.06225]	[0.00944, 0.02486]
Type of infl. surpr.	Point	Mean	Point	Mean
R^2	0.008	0.002	0.008	0.001
Obs.	101,438	92,765	101,438	92,765

(c) 1 year inflation, squared changes

	Dep. var.: $(\Delta m_{i,t})^2$			
	Contemp. uncertainty		Lagged uncertainty	
	(1)	(2)	(3)	(4)
$(\text{Unc.})^2$	0.00078***	0.00040***	0.00084***	0.00042***
	[0.00068, 0.00088]	[0.00034, 0.00046]	[0.00074, 0.00095]	[0.00037, 0.00047]
Type of infl. expectation	Point	Mean	Point	Mean
R^2	0.012	0.008	0.017	0.010
Obs.	102,140	116,720	102,140	116,720

(d) 3 year inflation, squared changes

	Dep. var.: $(\Delta m_{i,t})^2$			
	Contemp. uncertainty		Lagged uncertainty	
	(1)	(2)	(3)	(4)
$(\text{Unc.})^2$	0.00071***	0.00022***	0.00081***	0.00030***
	[0.00061, 0.00081]	[0.00017, 0.00027]	[0.00071, 0.00091]	[0.00025, 0.00035]
Type of infl. expectation	Point	Mean	Point	Mean
R^2	0.007	0.002	0.010	0.005
Obs.	101,059	117,019	101,059	117,019

Note: OLS regressions of changes in bin-mean expectations $\Delta m_{i,t}$ on the inflation surprise $\hat{\pi}_{i,t}$ and on the interaction of the surprise with subjective uncertainty $u_{i,t}\hat{\pi}_{i,t}$. Columns (1)–(2) use contemporaneous uncertainty $u_{i,t}$; columns (3)–(4) use lagged uncertainty $u_{i,t-1}$. Within each pair, the first column constructs the inflation surprise from the respondent's point forecast (lagged one month) and the second from the distributional mean (lagged two months). Panel (a) reports results for 1-year expectations and panel (b) for 3-year expectations. Panels (c) and (d) report OLS regressions of the squared change in expectations $(\Delta m_{i,t})^2$ on squared uncertainty, contemporaneous $u_{i,t}^2$ in columns (1)–(2) and lagged $u_{i,t-1}^2$ in columns (3)–(4); within each pair, the dependent variable is constructed from the point forecast (Point) or the distributional mean (Mean), with panel (c) using 1-year and panel (d) 3-year expectations. These regressions are the OLS analogues, without the control function, of the second-stage estimates in table 5. 95% confidence intervals based on robust standard errors allowing for two-way clustering by respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.2 First stage

The top panel of figure 3 illustrates the first stage graphically, plotting average uncertainty by tenure in the survey. There is a clear negative relationship for both 1- and 3-year horizons, and it is statistically well estimated. Much of the decline in average uncertainty occurs in the first few months of tenure, but even between months 7 and 12 there is a decline in uncertainty of about 14 percent. In principle the decline could be caused by selective attrition from the sample, but figure A.1 shows that the decline is the same restricting to the subset of respondents who stay in the sample for all 12 months (see section 4.5).

Table 2 reports estimation results for the baseline linear first stage (equation (11)). For all results here and in the rest of the paper, standard errors are heteroskedasticity robust with two-way clustering by person and time. The top panel reports results for one-year expectations and the bottom panel three-year expectations. The first two columns in each panel report results for regressions of uncertainty on tenure, restricting to either tenure>1 month or tenure>2 months, which corresponds to the two different subsamples that are used depending on the lag used for expectations in calculating inflation surprises.

In columns 1 and 2 in both panels, tenure is highly statistically significant, and the F-statistics are greater than 100, implying there are no concerns about weak instruments in this case.

Columns 3 and 4 control for inflation surprises in addition to tenure. They differ depending on whether inflation surprises are measured based on the distributional mean or point forecast. In that case the inflation surprises absorb some of the effect of tenure because expectations themselves – which are part of the inflation surprise calculation – are also correlated with tenure. Tenure, though, remains strongly related to uncertainty.

While the baseline specification is a typical linear first stage, the tenure-uncertainty relationship is fairly clearly nonlinear in the data. Figure A.3 and table A.11 report first-stage results for an alternative setup in which the right-hand side involves 1/tenure as a way of modeling that nonlinearity and shows that it fits much better. Second-stage results in this case are reported below and are highly similar to the baseline.

Table A.1 in the appendix reports results for the 2SLS first stage, which is a regression of the inflation surprise multiplied by uncertainty on the surprise and tenure times the surprise (i.e. it is the CF first stage with all variables interacted with the inflation surprise). While it is true that the F-statistics are large in all cases, they are driven by the presence of the inflation surprise on the right-hand side. The robust partial F-statistic on the excluded instrument – $\hat{\pi}_{i,t}z_{i,t}$ – is equal here to the squared t-statistic, which is 38.1 and 15.3 in the two columns of panel (a), and 18.8 and 20 for panel (b). The Montiel-Olea–Pflueger (2013)

Table 2: First stage regressions

(a) First Stage, 1 year inflation

	(1)	(2)	(3)	(4)
Dep var:	Unc.	Unc.	Unc.	Unc.
Tenure	−0.377***	−0.264***	−0.183***	−0.195***
	[−0.421, −0.333]	[−0.310, −0.218]	[−0.221, −0.145]	[−0.241, −0.149]
Infl. Surpr.			−0.602***	0.177***
			[−0.757, −0.447]	[0.078, 0.276]
Type of infl. surpr.			Point	Mean
F-statistic	279.182	124.769	87.305	38.462
R^2	0.004	0.002	0.029	0.003
Observations	116,720	104,595	103,888	93,460

(b) First Stage, 3 year inflation

	(1)	(2)	(3)	(4)
Dep var:	Unc.	Unc.	Unc.	Unc.
Tenure	−0.393***	−0.294***	−0.183***	−0.231***
	[−0.438, −0.348]	[−0.341, −0.247]	[−0.223, −0.143]	[−0.277, −0.185]
Infl. Surpr.			−0.601***	0.271***
			[−0.755, −0.447]	[0.176, 0.366]
Type of infl. surpr.			Point	Mean
F-statistic	294.550	151.649	89.453	59.977
R^2	0.004	0.002	0.026	0.005
Observations	117,019	104,827	101,438	92,765

Note: The table reports results of regressions of uncertainty on tenure and inflation surprises, with 1-year inflation expectations in panel (a) and 3-year expectations in panel (b). The first column restricts the sample to those with tenure above 1, the second to those with tenure above 2. Columns (3) and (4) use CPI headline inflation surprises constructed from the point forecast (lagged one month) and the bin mean (lagged two months), respectively. The regressions correspond to the first stage of the control-function estimation approach. 95% confidence intervals based on robust standard errors allowing for two-way clustering by person and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

cutoff for this case (at their $x = 10$) is 23.1, indicating that 2SLS is weakly identified here, demonstrating an important advantage of the CF approach (see Wooldridge (2015)).

4.3 Reduced-form estimates

4.3.1 News regressions

Columns (1) and (2) of table 3 report results for reduced-form regressions of $\Delta m_{i,t}$ on $\hat{\pi}_{i,t}$ and $\hat{\pi}_{i,t} z_{i,t}$. That is, compared to table 1, uncertainty is replaced with tenure. Since tenure is negatively associated with uncertainty, if it is also exogenous to attention in the cross-section

the model predicts that the coefficient on the interaction should be negative. That is in fact what we observe across all the regressions in table 3.

Table 3: Reduced form regressions

(a) 1 year inflation

	Dep. var.: Δm			Dep. var.: $(\Delta m)^2$
	(1)	(2)		(3)
Tenure × Infl. Surpr.	−0.0046***	−0.0021***	Tenure	−0.26***
	[−0.0057, −0.0035]	[−0.0032, −0.0010]		[−0.28, −0.24]
Infl. Surpr.	0.1112***	0.0445***		
	[0.0931, 0.1293]	[0.0326, 0.0564]		
Type of infl. surpr.	Point	Mean		
R^2	0.0270	0.0040	R^2	0.01
Obs.	103,888	93,460	Obs.	116,720

(b) 3 year inflation

	Dep. var.: Δm			Dep. var.: $(\Delta m)^2$
	(1)	(2)		(3)
Tenure × Infl. Surpr.	−0.0025***	−0.0007	Tenure	−0.27***
	[−0.0036, −0.0014]	[−0.0018, 0.0004]		[−0.29, −0.25]
Infl. Surpr.	0.0611***	0.0244***		
	[0.0472, 0.0750]	[0.0126, 0.0362]		
Type of infl. surpr.	Point	Mean		
R^2	0.0070	0.0010	R^2	0.01
Obs.	101,438	92,765	Obs.	117,019

Note: The left side of the table shows reduced-form regressions of changes in bin mean expectations ($\Delta m_{i,t}$) on tenure interacted with inflation surprises, as well as the level of the inflation surprise. The right side regresses the squared changes in expectations $(\Delta m_{i,t})^2$ on tenure. Columns (1)–(2) vary the construction of the inflation surprise using either the point forecast (lagged by one month) or the bin mean (lagged by two months). 95% confidence intervals based on robust standard errors allowing for two-way clustering by respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To quantify magnitudes, the average sensitivity of one-year expectations to news at $tenure = 2$ in columns (1) and (2) is 0.10 and 0.04, compared to 0.06 and 0.02 at 12 months respectively for the two specifications. A person at the end of their time in the sample is therefore on average only half as responsive to inflation surprises as at the beginning. The sensitivity of three-year expectations also declines by about half. Table A.2 shows that similar results hold across the range of permutations of the specification.

The middle panel of figure 3 plots estimates from regressions of $\Delta m_{i,t}$ on tenure dummies multiplied by $\hat{\pi}_{i,t}$, yielding estimates of sensitivity to surprises at each level of tenure without

imposing linearity in the interaction. The negative relationship is clearly apparent. Again, the estimates imply that sensitivity to inflation surprises falls by about half for agents at the end of their tenure compared to the beginning and that is consistent with the result from the top panel that uncertainty also falls by about half over the same period.

Note, importantly, that the decline by a factor of one half in both uncertainty and responsiveness is consistent with the model's prediction that sensitivity to news is proportional to uncertainty.

Kim and Binder (2023) show that in addition to uncertainty declining with tenure, average inflation expectations do, also. That should cause the average of $\Delta m_{i,t}$ to be negative. If the rate of decline changes with tenure, then tenure would naturally also be included as a control in these regressions. Table A.3 in the appendix reports results for regressions analogous to those in table 3 but including tenure as an additional control. The coefficient on the interaction $\hat{\pi}_{i,t} z_{i,t}$ is essentially unaffected. The estimated coefficient on tenure is positive and the constant in the regression negative, consistent with average expectations declining, but at a slowing rate, with tenure.

4.3.2 Volatility in expectations

The third column of table 3 reports results from regressions of $(\Delta m_{i,t})^2$ on tenure. The coefficient is again significantly negative, implying that as tenure increases agents are less responsive to information overall (under the model that is because their expectations become more stable). The bottom panel of figure 3 plots the average of $(\Delta m_{i,t})^2$ by tenure bins. The negative relationship holds across the entire range of tenure and is not driven by any particular bin. Note also that the decline is very similar in shape to the decline in uncertainty itself. Again, consistent with the proportionality prediction of the model, the approximately 50% proportional decline in belief volatility matches the proportional decline in uncertainty across tenure bins.

4.4 Second-stage estimates

4.4.1 News regressions

Table 4 reports results from the second stage of the control function estimator – equation (15) – where the conditional expectations are, in the benchmark results, approximated as linear functions of $\eta_{i,t}$ and $\hat{\pi}_{i,t}$.[13] The first row reports the mean sensitivity of expectations to $u_{i,t}\hat{\pi}_{i,t}$

[13]Standard errors here are constructed by putting the first and second stage regressions into a single GMM problem so that the standard errors reported in table 4 account for uncertainty in the first stage. As above,

(i.e. at the average value of the various interactions) which, under the model, represents the cross-sectional mean of the precision of the signal each agent gets from inflation, $\sigma_{i,t}^{-2}$. (The full specification of the regression is reported in appendix table A.4.) That coefficient ranges between 0.0032 and 0.0229, depending on the specification, and it is statistically well estimated (insignificant only in column (4)). The point estimates imply that the standard deviation in the noise in agents' signals about future inflation from current realized inflation is between 7 and 18 percent. Recall from the model that that noise combines both the gap between realized inflation and its persistent component and also the difference between what agents observe and actual realized inflation.

Table 4: Second Stage Regression

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr., AME	0.0229***	0.0114***	0.0120***	0.0032
	[0.0163, 0.0295]	[0.0049, 0.0179]	[0.0057, 0.0183]	[−0.0019, 0.0083]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0245***	−0.0127***	−0.0134***	−0.0046*
	[−0.0311, −0.0179]	[−0.0191, −0.0063]	[−0.0198, −0.0070]	[−0.0098, 0.0006]
Other controls	Y	Y	Y	Y
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	103,888	93,460	101,438	92,765

Note: Results from second-stage regressions of $\Delta m_{i,t}$ (measured based on the bin mean) on interactions of uncertainty (Unc. = $u_{i,t}$), inflation surprise ($\hat{\pi}_{i,t}$), and control function ($\eta_{i,t}$). The regression includes the terms obtained from the Kroneker product of (inflation surprise; uncertainty×inflation surprise) with (a constant, the control function component, and inflation surprise). All the coefficients of this regression are reported in table A.4. AME denotes average marginal effect. This table reports the mean sensitivity of expectations to $u_{i,t}\hat{\pi}_{i,t}$ in the first row and the mean sensitivity of expectations to $\eta_{i,t}\hat{\pi}_{i,t}$ (both computed at the average values of the various interactions). Columns (1)–(2) use different versions of the inflation surprise—measured based on the point forecast (Point) or the lagged bin mean (Mean)—for the 1-year inflation expectations. Columns (3)–(4) report the analogous specifications for the 3-year inflation expectations. 95% confidence intervals in brackets account for sampling uncertainty in the first stage and allow for clustering by both respondent and time. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The second row of table 4 reports the mean sensitivity of expectations to $\eta_{i,t}\hat{\pi}_{i,t}$, which measures how the sensitivity of expectations to news depends on the *endogenous* part of uncertainty – the part not driven by tenure. The fact that the coefficients are negative is consistent with the idea that agents with high uncertainty are probably people who pay less attention to inflation on average, so that $\sigma_{i,t}^{-2}$ is low, with the result that the net relationship between sensitivity to inflation and $\eta_{i,t}$ need not be positive.

Table A.5 reports results from the other 12 permutations of the specification choices (CPI vs PCE, etc.). While there is variation in the coefficient estimates for $u_{i,t}\hat{\pi}_{i,t}$, they

standard errors are two-way clustered by person and time.

are broadly consistent with each other and uniformly positive and significant. Table A.6 replicates table 4 but using agents' point forecast for inflation on the left-hand side instead of their implied mean. Finally, table A.7 reports results from second-stage regressions including richer controls – most importantly, higher powers of various variables (recalling the sieve interpretation of the regression) – and finds similar results.

For the sake of completeness, table A.10 in the appendix reports the 2SLS second stage estimates. The coefficients are similar to those obtained via the CF method, as are the standard errors, suggesting that in practice both the cross-sectional heterogeneity and first-stage concerns with 2SLS may not be particularly important.

4.4.2 Volatility in expectations

Table 5 reports results for the second-stage regression of $(\Delta m_{i,t})^2$ on uncertainty – equation (16). The estimated sensitivity to $u_{i,t}^2$ is, under the conditions discussed above, again an estimate of $\mathbb{E}\left[\sigma_{i,t}^{-2}\right]$, where here that represents the precision of the total signal that the agents receive in period t. Intuitively, the idea behind the regression is that, all else equal, when an agent's uncertainty is higher, the variance of the change in their expectations should also be higher, in proportion to the precision of their signals.

Table 5: Second Stage Regression for $(\Delta m)^2$

	1-year inflation		3-year inflation	
	(1)	(2)	(3)	(4)
$(Unc.)^2$, AME	0.1336***	0.0763***	0.1902***	0.0753***
	[0.1038, 0.1634]	[0.0661, 0.0864]	[0.1493, 0.2311]	[0.0655, 0.0852]
$\eta_{i,t}^2$, AME	0.0099*	0.0096***	0.0338***	0.0167***
	[−0.0016, 0.0213]	[0.0061, 0.0132]	[0.0196, 0.0480]	[0.0136, 0.0198]
Type of inflation expectation	Point	Mean	Point	Mean
Observations	102,140	116,720	101,059	117,019

Note: Results from second-stage regressions of $(\Delta m_{i,t})^2$ on interactions of uncertainty and the control-function component. Uncertainty is $Unc_{i,t} = u_{i,t}$. The regression includes terms obtained from the Kronecker product $(1, Unc_{i,t}^2) \otimes (1, \eta_{i,t}, \eta_{i,t}^2, \eta_{i,t}^{\text{inv}})$, where $\eta_{i,t}$ is demeaned and $\eta_{i,t}^{\text{inv}} = (\eta_{i,t} - \min(\eta_{i,t}) + 10)^{-1}$. All coefficients from the second-stage regression are reported in the appendix. This table reports the AME (average marginal effect) of $Unc_{i,t}^2$ and the AME of the control-function component, both evaluated at the sample-average values of the interaction terms. Columns (1)–(2) use 1-year inflation expectations and columns (3)–(4) use 3-year inflation expectations; within each horizon, the dependent variable is constructed using point expectations (Point) or lagged bin means (Mean). 95% confidence intervals in brackets account for sampling uncertainty in the first stage and allow for clustering by both respondent and time. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The estimated values of $\mathbb{E}\left[\sigma_{i,t}^{-2}\right]$ are significantly larger in table 5 than in table 4. That is exactly what we should expect – agents receive many signals about future inflation beyond

realized current inflation. The results in table 4 indicated that the signal agents got from inflation had an error with effective standard deviation of 7–18%. Table 5 implies that the standard deviation of the effective error in their *overall* signal is between 2.3% and 3.6%.

The second row of table 5 reports the sensitivity to $\eta_{i,t}^2$. In this case, in contrast to table 4, the coefficients are positive. The intuition for that result is as follows. Suppose some agents just pay less attention to the economy and know less about inflation, thus reporting distributions that tend to be wider. In table 4, the negative coefficient on $\eta_{i,t}\hat{\pi}_{i,t}$ indicates those agents' inflation expectations are less responsive to realized inflation than average. At the same time, it is plausible that those agents also report expectations that vary somewhat erratically – given that they do not pay much attention, there is not much reason to expect any real stability in their reported beliefs. That is why the coefficient in the second row of table 5 is positive.

It should be noted that that suggests a potential source of bias in the volatility regressions. Some of the variation in Δm_t may just be noise – e.g. due to difficulty people have in translating expectations into specific numbers. If that noise declines with tenure, that will create an upward bias in the estimate of $\mathbb{E}\left[\sigma_{i,t}^{-2}\right]$. The regressions for the trend-plus-noise model above are not exposed to this issue because any noise in expectations is absorbed into the residual. The next section addresses this issue further.

4.5 Robustness

This section examines a number of variations to the baseline analysis to address various possible concerns.

First, table A.2 reports reduced-form results for the various combinations of left- and right-hand side variables (CPI vs PCE, etc.). Table A.3 includes tenure as an additional control. Table A.12 reports second-stage results for the case where the first stage uses 1/tenure to account for the nonlinear relationship with uncertainty.

One might worry that the results are driven by unsophisticated respondents who report extreme and erratic beliefs. Figure A.2 and table A.9 in the appendix replicate the analysis for a subsample of more sophisticated respondents – those with a college degree or higher, household income of $75,000 or more (roughly the median US household income), and a correct answer to a numeracy question – for whom average uncertainty is less than half that of the full sample.[14] The results are highly similar to the baseline.

To address possible selective attrition in the sample, figure A.1 and table A.8 show that

[14]The survey includes four numeracy questions. We keep the one with the lowest number of correct answers (44%), which asks about compounding 10% annual interest over two years.

the results are similar for the subsample of respondents who complete all 12 interviews.

In all of the results for the trend-plus-noise model, the inflation surprise is measured based on CPI and PCE inflation in the same month as the survey. Since those numbers are not released until after the survey occurs, the official releases are not actually part of agents' information sets. Table A.13 reports second-stage results using announced inflation for the previous month and again obtains similar coefficient estimates.

The stability of the results for sensitivity to inflation might make one concerned that what is driving the results is not realized inflation itself but instead variation in inflation expectations. The first point to make would be that the model treats expectations and realizations symmetrically, in the sense that updates of expectations for agents with high uncertainty should be relatively weakly responsive both to realized inflation and to their own expectations. Past work, including Goldstein (2023) and Gemmi and Mihet (2026) uses mean reversion in expectations alone – without specifying the signal – to identify the strength of belief updating.

Table A.14 in the appendix reports three variations on the main analysis to try to isolate just the contribution of realized inflation. First, instead of using each agent's expectation from the previous period, rows 2–3 replace it with the cross-sectional mean or median expectation from the previous period. The results are similar to the baseline, but slightly weaker (as the model would imply, since we are no longer using the theoretically correct measure of expectations), and the results for 3-year inflation using the mean expectation are, as in the baseline, weakest (in fact with the wrong sign in this case).

Second, even weaker, rows 4–5 of the table report results just putting in the full-sample mean and median, respectively, of expectations – i.e. all the same constant for all agents in all periods, so that the surprise varies only through realized inflation. The results again point in the same direction but are weaker still. Using the full-sample mean, the estimates based on the point forecast remain positive and statistically significant at roughly one-half to two-thirds of the baseline, so uncertainty scales the response of expectations to variation coming from realized inflation alone; using the full-sample median, the estimates are smaller and insignificant. As in rows 2–3, the 3-year results based on the density mean are slightly negative.

Third, since expectations systematically decline with tenure (which is consistent with the model, since agents typically come into the sample with expectations that are too high), one might worry that the results are somehow driven by that decline. Row 6 of table A.14 therefore residualizes the expectations with respect to tenure dummies – i.e. removes a tenure trend – rebuilding both the change in expectations and the surprise from the residualized

expectations. The results are again similar to the baseline, but with smaller coefficients.

Again, a benefit of the analysis based on squared changes in expectations is that it does not require measuring the inflation surprise, so none of the issues addressed in table A.14 are relevant for those regressions. On the other hand, if there is any kind of measurement error in expectations whose magnitude varies systematically with tenure, that could bias this set of regressions. Table A.15 reports results for squared changes in expectations after a correction for noise described in appendix A.2 and after taking out a tenure effect, and finds results that are similar and in general not statistically significantly different from the baseline.

4.6 External validity

Recall that the paper's identifying assumption is that people pay more attention to inflation during the survey than they did prior to the survey. That fundamentally means that the people in the survey are different from the average person, and in particular that the estimated signal precision is higher than what would be expected for the typical person outside the survey. The results should therefore be viewed as giving an overestimate of or upper bound on the signal precision of the average person in the economy.

On the other hand, the paper's objective is not just to measure average signal precision. It is also to test whether people can be reasonably described as updating according to Bayes' formula in the sense that the sensitivity of their beliefs to information depends on uncertainty multiplied by signal precision. The empirical results here provide evidence in support of that view. The paper's identifying assumption that people in the survey have higher precision signals does not also imply that they are somehow "more Bayesian" than people outside the survey. So the finding that uncertainty determines how strongly expectations are anchored has a relatively stronger claim to external validity.

5 Time-series analysis

While the panel analysis is useful for directly testing models of beliefs, in practice policymakers typically focus on the cross-sectional mean or median of beliefs. Additionally, what is relevant for policymakers is most likely *forecasting* the sensitivity of beliefs to signals, as opposed to the contemporaneous relationships the analysis in the previous section tests. This section therefore examines whether current reported uncertainty in the SCE predicts either future sensitivity of average expectations to inflation or the magnitude of squared changes in beliefs. It also examines whether the aggregate relationship is consistent with the structural estimate

of the average of $\sigma_{i,t}^{-2}$.

5.1 The sensitivity of mean expectations to realized inflation

The paper's core objective is to understand the sensitivity of expectations to signals. Clearly at any given time, the average sensitivity is the average of $u_{i,t}\sigma_{i,t}^{-2}$. To see why that matters, consider a shock that raises uncertainty of all agents by some amount Δ_u, leaving signal precision fixed:

$$\frac{\partial}{\partial \Delta_u}\frac{\partial Avg_t\left(m_{i,t}\right)}{\partial \pi_t} = Avg_t\left(\sigma_{i,t}^{-2}\right) \tag{19}$$

where $Avg_t\left(\cdot\right)$ denotes the cross-sectional average on date t.[15] In words, an exogenous shift in uncertainty raises the sensitivity of expectations to realized inflation – and, more generally, to any signal – in proportion to the average signal precision. That average is what the second-stage regressions above identify. Equation (19) shows formally that really we are after a second-order sensitivity – the response of the sensitivity of expectations to a change in uncertainty.

This section, instead of using instrumental variables, just uses time-series variation. Specifically, if everything is collapsed into cross-sectional averages and covariances, we have, from equation (9),

$$\begin{aligned}\Delta Avg_t\left[m_{i,t}\right] &= Avg_t\left[u_{i,t}\right]\left(\pi_t - Avg_t\left[m_{i,t-1}\right]\right)Avg_t\left[\sigma_{i,t}^{-2}\right] \\ &\quad -Cov_t\left[u_{i,t}, m_{i,t-1}\right]Avg_t\left[\sigma_{i,t}^{-2}\right] + Cov_t\left[u_{i,t}\left(\pi_t - m_{i,t-1}\right), \sigma_{i,t}^{-2}\right] + Avg_t\left[\epsilon_{i,t}\right]\end{aligned} \tag{20}$$

The terms in the first line are observable, leaving $Avg_t\left[\sigma_{i,t}^{-2}\right]$ as a coefficient to be estimated. We treat it as a constant for the sake of the regressions. The second line includes various terms that are unobservable. To the extent that they are constant over time, or uncorrelated with $Avg_t\left[u_{i,t}\right]\left(\pi_t - Avg_t\left[m_{i,t-1}\right]\right)$, the regression will be identified. Certainly they may not be constant in reality; the cross-sectional identification above is the tighest identification the paper has. Instead, equation (20) just shows what the model implies for a standard time-series regression.

In the cross-section one of the major sources of endogeneity was the correlation between $\sigma_{i,t}^{-2}$ and $u_{i,t}$ – people who have low attention have high uncertainty. The time-series regression will be identified if that effect is *only* in the cross-section, i.e. if time-series variation in uncertainty is unrelated to time-series variation in signal precision.

[15] Equation (19) follows directly from results in Dytso, Poor and Shitz (2022) and is exact, rather than being a derivative of the approximation (9).

For squared changes in expectations, the analogous equation is

$$Avg_t\left[(\Delta m_t)^2\right] \approx Avg_t\left[\left(E_{t-1}\left[u_{i,t}\right]\right)^2\right] Avg_t\left[\sigma_{i,t}^{-2}\right] + Cov_t\left[\left(E_{t-1}\left[u_t\right]\right)^2, \sigma_t^{-2}\right] \tag{21}$$

5.2 Estimates

To align with typical aggregate time series, we first collapse the panel structure into cross-sectional means and medians in each period (medians are included because policymakers often focus on them). We then define inflation news as inflation in month t minus the cross-sectional mean or median of 12-month expected inflation (from the distributional mean) in month $t-2$. The expectation is lagged by two months because the change in inflation, which is on the left-hand side of the regression, involves expected inflation in month $t-1$.

Table 6 reports results from eight regressions. The first four are of the change in mean and median expectations at the one- and three-year horizons on inflation news, both alone and interacted with the cross-sectional median of uncertainty at date $t-1$ (u_{t-1}).

Table 6: Aggregate Analysis

	(1)	(2)	(3)	(4)
Dep. var.:	Δm (1yr)	Δm (3yr)	$(\Delta m)^2$ (1yr)	$(\Delta m)^2$ (3yr)
Infl. Surprise	−0.00563	−0.00117		
	[−0.02815, 0.01689]	[−0.01755, 0.01521]		
Lagged unc.	−0.01638	−0.01846		
	[−0.04786, 0.01510]	[−0.04673, 0.00981]		
Surprise × Lagged unc.	0.00870***	0.00456**		
	[0.00371, 0.01369]	[0.00028, 0.00884]		
Lagged (unc.)^2			0.00372***	0.00197***
			[0.00229, 0.00515]	[0.00093, 0.00301]
Constant	0.05455	0.05497	−0.00307	0.00018
	[−0.01769, 0.12679]	[−0.01883, 0.12877]	[−0.01464, 0.00850]	[−0.00864, 0.00900]
F-statistic	8.62	13.24	25.89	13.82
R^2	0.223	0.120	0.262	0.138
Observations	140	140	141	141

Note: Column (1) reports results of a time-series regression of the change in aggregate inflation expectations (cross-sectional median of individual expectations) on CPI inflation surprise, computed as realized inflation minus two-month-lagged cross-sectional median expectation, lagged aggregate uncertainty (cross-sectional median of individual uncertainty), and their interaction. Column (2) repeats the exercise using 3-year inflation expectations and uncertainty. Column (3) regresses squared changes in 1-year aggregate inflation expectations on lagged squared uncertainty. Column (4) repeats the exercise using 3-year changes in inflation expectation and uncertainty. 95% confidence intervals based on Newey-West standard errors with 12 lags are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In the sample, the standard deviations of one- and three-year uncertainty are 1.4 and 1.0,

respectively. In the first column, the coefficient on the interaction of news with uncertainty is 0.009, while the coefficient on news alone is actually slightly negative and not statistically distinguishable from zero. Both of those are consistent with the model. The regression implies that if, hypothetically, people had zero uncertainty about future inflation, their expectations would be completely unresponsive to realized inflation. Additionally, the value of the coefficient on uncertainty is consistent with the structural estimate of mean signal precision from table 4. That is, the variation in the sensitivity of average expectations that is related to variation in average uncertainty is approximately of the magnitude that would be expected from the variation in sensitivity observed across months of tenure within respondents.

The results are similar for three-year expectations. Again, the coefficient on the inflation surprise by itself is statistically indistinguishable from zero, and the coefficient on the interaction with uncertainty is 0.005, a similar magnitude to the micro estimate.

The left-hand panel of figure 4 plots the fitted sensitivity of one- and three-year expectations from the two regressions. There is a striking degree of variation. At its peak in 2022, the sensitivity of one-year expectations is estimated to be 0.06. Exactly when inflation was highest, expectations were most sensitive to inflation surprises. That sensitivity was higher than the value at the beginning of 2020 by a factor of four. Furthermore, at the end of the sample in 2025, sensitivity remains nearly twice as high as it was in 2020.

Similar results, though at lower overall levels, hold for three-year expectations. Their sensitivity to inflation surprises peaked at 0.04, and they are also now about twice their pre-covid values.

The second pair of columns in table 6 report results from regressions of squared changes in median expectations on lagged squared uncertainty. As in the first two columns, the constants in the regressions are not statistically significant, consistent with the model.

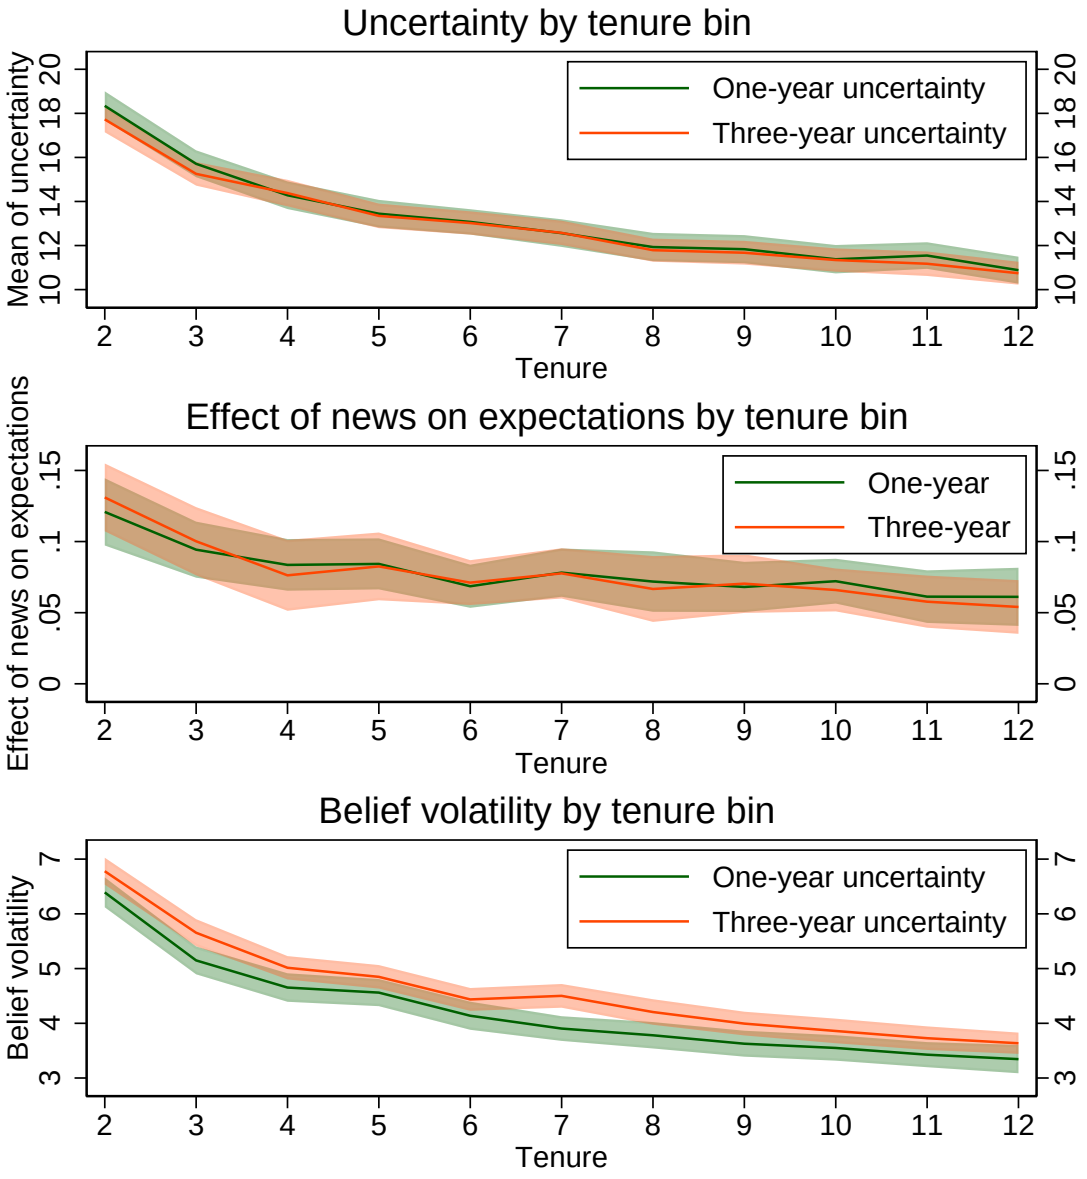
The coefficients on uncertainty are again highly statistically significant, but in this case, unlike the first two columns, the estimated coefficients do not line up well with those from the person-level estimates. That is consistent with the potential bias discussed above if there is noise in individuals' reports of their expectations.

Again, the coefficients on squared uncertainty in these columns should be estimates of average signal precision. They are smaller than those in the first two columns, but of broadly similar magnitude, and their confidence bands overlap substantially.

To see the effects of time-series variation in uncertainty on the predicted volatility, the right-hand panel of figure 4 plots the square roots of the fitted values from the regressions from columns (3) and (4) of table 6. These values represent conditional standard deviations

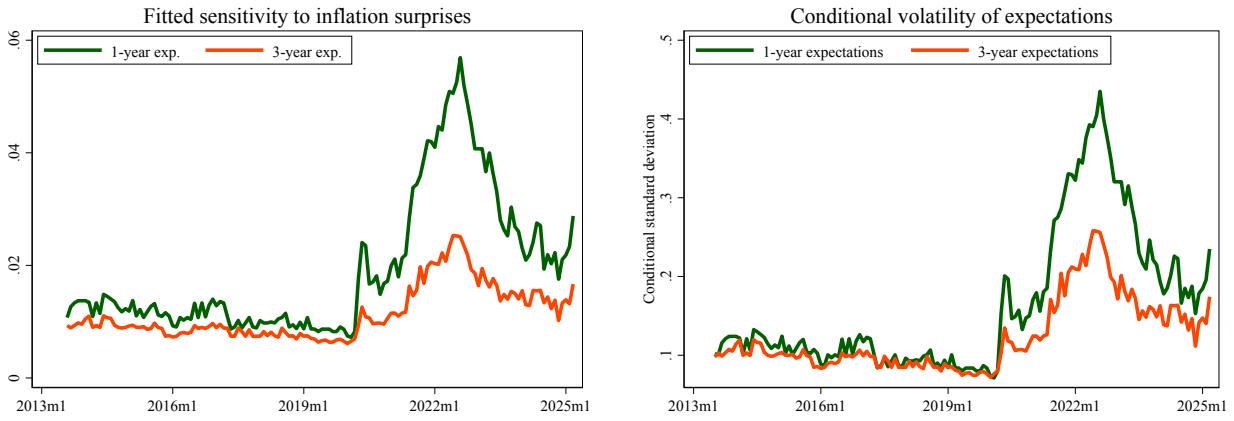
for one- and three-year median inflation expectations under the assumption that the expectations are martingales. Those standard deviations were both equal to about 0.07 at the beginning of 2020. They rose by a factor of 4–6 in 2022 and by 2025 remained more than twice as high as pre-covid.

Figure 3: First stage and reduced-form relationships



Note: In each panel, the lines are point estimates and the shaded regions 95-percent confidence bands. Standard errors are two-way clustered by person and time. The top panel reports coefficients from a regression of uncertainty (each agent's variance) on tenure indicators. The middle panel is from a regression of the change in each agent's conditional mean on inflation news interacted with tenure indicators. The bottom panel is from a regression of the squared change in expectations on tenure indicators.

Figure 4: Aggregate time-series regressions: sensitivity and conditional volatility



(a) Fitted sensitivity to inflation surprises

(b) Fitted conditional volatility of expectations

Note: Panel (a) plots fitted values of the time-varying sensitivity of changes in median inflation expectations to inflation surprises for the 1-year and 3-year horizons. Panel (b) plots the square roots of fitted values from the regressions of squared changes in median expectations on lagged squared uncertainty. Standard errors in the underlying regressions are Newey–West with 12 lags.

6 Conclusion

This paper studies the anchoring of inflation expectations. Starting from the observation that the response function of expectations to news is fundamentally unobservable, its key insight is that under the assumptions that agents are Bayesian and that they observe Gaussian signals, the response function can be recovered from knowledge of agents' posterior distributions. And in fact the recovery is not even particularly complicated: the derivatives of the response function are simply the posterior cumulants.

Obviously the structural assumptions are strong. That is somewhat inevitable since the goal is to measure something unobservable. Ultimately the empirical results have to speak for themselves: does an agent's reported uncertainty actually have any predictive power for the sensitivity of their beliefs to news? Across a variety of measures – contemporaneous and forward-looking, time-series and cross-sectional – the answer is broadly yes.

The results imply that the period of high inflation in 2021 and 2022 has had, so far, lasting effects on the strength of the expectational anchor. Agents' uncertainty about future inflation rose as inflation did, and subsequently declined, but not nearly to the extent that inflation itself did. By the beginning of 2025 inflation was approaching its pre-2020 levels, but inflation uncertainty was still about two times higher. Those facts show how the results in this paper can be used going forward. They provide a real-time quantitative measure of how well inflation expectations are anchored. More generally, they show how to measure the sensitivity of expectations to news in arbitrary settings.

References

Ameriks, John, Andrew Caplin, Steven Laufer, and Stijn Van Nieuwerburgh, "The joy of giving or assisted living? Using strategic surveys to separate public care aversion from bequest motives," *The journal of finance*, 2011, *66* (2), 519–561.

Armantier, Olivier, Argia Sbordone, Giorgio Topa, Wilbert van der Klaauw, and John C. Williams, "A New Approach to Assess Inflation Expectations Anchoring Using Strategic Surveys," *Journal of Monetary Economics*, 2022, *129*, S82–S101.

_ , Scott Nelson, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar, "The price is right: Updating inflation expectations in a randomized price information experiment," *Review of Economics and Statistics*, 2016, *98* (3), 503–523.

_ , **Wändi Bruine de Bruin, Giorgio Topa, Wilbert Van Der Klaauw, and Basit Zafar**, "Inflation expectations and behavior: Do survey respondents act on their beliefs?," *International Economic Review*, 2015, *56* (2), 505–536.

Ball, Laurence M. and Sandeep Mazumder, "Inflation Dynamics and the Great Recession," *Brookings Papers on Economic Activity*, 2011, pp. 337–381.

Binder, Carola Conces, Jeffrey R. Campbell, and Jane M. Ryngaert, "Consumer Inflation Expectations: Daily Dynamics," 2024.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, "Diagnostic expectations and stock returns," *The Journal of Finance*, 2019, *74* (6), 2839–2874.

Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston, "Anchored inflation expectations," *American Economic Journal: Macroeconomics*, 2023, *15* (1), 1–47.

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia, "Inflation expectations, learning, and supermarket prices: Evidence from survey experiments," *American Economic Journal: Macroeconomics*, 2017, *9* (3), 1–35.

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

_ **and** _ , "Inflation Expectations and Monetary Policy?What Have We Learned and To What End?," 2025. Working paper.

D'Acunto, Francesco and Michael Weber, "Why Survey-Based Subjective Expectations areMeaningful and Important," *Annual Review of Economics*, 2024.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance*, 1998, *53* (6), 1839–1885.

DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly journal of economics*, 2003, *118* (3), 909–968.

Dew-Becker, Ian, Stefano Giglio, and Pooya Molavi, "Learning and the emergence of nonlinearity in financial markets," 2026. Working paper.

Dytso, Alex, H Vincent Poor, and Shlomo Shamai Shitz, "Conditional mean estimation in Gaussian noise: A meta derivative identity with applications," *IEEE Transactions on Information Theory*, 2022, *69* (3), 1883–1898.

Engle, Robert F, "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation," *Econometrica: Journal of the Econometric Society*, 1982, pp. 987–1007.

Farmer, Leland E, Emi Nakamura, and Jón Steinsson, "Learning about the long run," *Journal of Political Economy*, 2024, *132* (10), 000–000.

Fleckenstein, Matthias, Francis A Longstaff, and Hanno Lustig, "Deflation risk," *The Review of Financial Studies*, 2017, *30* (8), 2719–2760.

Gemmi, Luca and Roxana Mihet, "Household Expectations in Uncertain Times," 2026. Working paper.

Goldstein, Nathan, "Tracking inattention," *Journal of the European Economic Association*, 2023, *21* (6), 2682–2725.

Hall, Peter and Christopher C Heyde, *Martingale limit theory and its application*, Academic press, 2014.

Hilscher, Jens, Alon Raviv, and Ricardo Reis, "How likely is an inflation disaster?," *The Review of Financial Studies*, 2025, p. hhaf058.

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

Kitsul, Yuriy and Jonathan H Wright, "The economics of options-implied inflation probability density functions," *Journal of Financial Economics*, 2013, *110* (3), 696–711.

Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko, "Inflation targeting does not anchor inflation expectations: Evidence from firms in New Zealand," *Brookings Papers on Economic Activity*, 2015, *2015* (2), 151–225.

Lamla, Michael J and Dmitri V Vinogradov, "Central bank announcements: Big news for little people?," *Journal of Monetary Economics*, 2019, *108*, 21–38.

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

Masten, Matthew A and Alexander Torgovitsky, "Identification of instrumental variable correlated random coefficients models," *Review of Economics and Statistics*, 2016, *98* (5), 1001–1005.

Meghir, Costas and Luigi Pistaferri, "Income Variance Dynamics and Heterogeneity," *Econometrica*, 2004, *72* (1), 1–32.

Mertens, Thomas M and John C Williams, "What to expect from the lower bound on interest rates: Evidence from derivatives prices," *American Economic Review*, 2021, *111* (8), 2473–2505.

Molavi, Pooya, "Misspecified Bayesianism," 2025. Working paper.

Newey, Whitney and Sami Stouli, "Control variables, discrete instruments, and identification of structural functions," *Journal of Econometrics*, 2021, *222* (1), 73–88.

Newey, Whitney K, "Convergence rates and asymptotic normality for series estimators," *Journal of econometrics*, 1997, *79* (1), 147–168.

Reis, Ricardo, "Why did inflation rise and fall in 2021-24?Channels and evidence from expectations," 2025. Working paper.

Roth, Christopher and Johannes Wohlfart, "How do expectations about the macroeconomy affect personal expectations and behavior?," *Review of Economics and Statistics*, 2020, *102* (4), 731–748.

Stock, James H and Mark W Watson, "Why has US inflation become harder to forecast?," *Journal of Money, Credit and banking*, 2007, *39*, 3–33.

Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos et al., "Tell Me Something I Don't Already Know: Learning in Low-and High-Inflation Settings," *Econometrica*, 2025, *93* (1), 229–264.

Wooldridge, Jeffrey M, "Control function methods in applied econometrics," 2015, *50* (2), 420–445.

A Additional results

A.1 Estimating precision by months of tenure

We use here a simplified version of equation (5) to get a rough estimate of signal precision at different levels of tenure. Starting from equation (5), first shut off the term involving the effect of signals – treating it as mean-zero. Second, denoting by c the contribution of uncertainty coming from innovations to π_t^* – and therefore implicitly taking it as constant – we have

$$\Delta u_t \approx c - \sigma_t^{-2} u_{t-1}^2 \tag{22}$$

So then in a steady-state prior to entering the survey, it must be that $c = \sigma_t^{-2} u_{t-1}^2$.

Now denote average uncertainty at tenure of n months by n. Assuming signal precision and uncertainty in month 1 are equal to their steady-state values, the change in average uncertainty at tenure of n months is

$$\Delta u_n \approx \sigma_1^{-2} u_1^2 - \sigma_n^{-2} u_{n-1}^2 \tag{23}$$

and so

$$\sigma_n^{-2} \approx -\frac{\Delta u_n - \sigma_1^{-2} u_1^2}{u_{n-1}^2} \tag{24}$$

In that equation, the only unknown term on the right-hand side is σ_1^{-2}. However, given σ_1^{-2}, we can then get σ_n^{-2} for all $n > 1$. To find σ_1^{-2}, we search for the value such that the average value for $1 \leq n \leq 12$ of σ_n^{-2} implied by (24) matches the estimates in table 4, meaning we have two values of σ_n^{-2}, one for each for each forecast horizon.

Figure A.4 plots the four estimates of σ_n^{-2} across values of tenure for $n > 1$ (since those are the values actually used in the regression and need to be assumed to be unrelated to tenure). In both cases, one of the estimates implies that precision is stable across tenure. For the one-year estimates, the other value implies that precision is increasing, while for the three-year estimates the other value implies it is decreasing.

The estimated values of σ_n^{-2} here are dependent on adding extra structure to the uncertainty equation and imposing a steady-state assumption in month 1, so we do not present them as credible estimates of signal precision. Instead, they are just meant to give a rough test of the identifying assumption of constant signal precision, and the evidence suggests that the assumption is reasonable.

A.2 Correcting for noise in squared changes in expectations

Suppose that reported expectations contain transitory noise – for example, error in how respondents map their beliefs into their survey answers. Squared changes in reported expectations then overstate the volatility of true beliefs, and if the magnitude of the reporting noise varies systematically with tenure, it could bias the second-stage regressions for squared changes in expectations. The noise can be accounted for using a method from the literature on income dynamics (Meghir and Pistaferri, 2004).

Specifically, suppose we observe some $y_{i,t}$ (not the signal from the paper; just arbitrary notation), with

$$y_{i,t} = x_{i,} + u_{i,t} \tag{25}$$

$$\Delta x_{i,t} = \varepsilon_{i,t} \tag{26}$$

where $x_{i,t}$ is a martingale and $u_{i,t}$ is serially uncorrelated noise. Then it is straightforward to confirm that

$$E\left[(\Delta y_{i,t})^2 + \Delta y_{i,t}\Delta y_{i,t-1} + \Delta y_{i,t}\Delta y_{i,t+1}\right] = \sigma^2_{\varepsilon,t} \tag{27}$$

The moment on the left-hand side therefore isolates the variance of the innovations to x, purging the contribution of the transitory noise. Table A.15 applies this correction to the dependent variable of table 5, replacing $(\Delta m_{i,t})^2$ with $\Delta m_{i,t}\left(\Delta m_{i,t-1} + \Delta m_{i,t} + \Delta m_{i,t+1}\right)$. The corrected coefficient estimates are somewhat weaker than the baseline, which is also consistent with there being variation in measurement error (communication error) across tenure.

Table A.1: First stage regressions, 2SLS approach

(a) First stage, 1 year inflation

	(1)	(2)
Dep var:	Unc. × Surp.	Unc. × Surp.
Tenure × Surpr.	−0.325***	−0.183***
	[−0.428, −0.222]	[−0.275, −0.091]
Infl. Surpr.	13.817***	12.228***
	[12.352, 15.282]	[11.243, 13.213]
Type of infl. surpr.	Point	Mean
F-statistic	188.379	319.883
R^2	0.217	0.229
Observations	103,888	93,460

(b) First stage, 3 year inflation

	(1)	(2)
Dep var:	Unc. × Surp.	Unc. × Surp.
Tenure × Surpr.	−0.266***	−0.226***
	[−0.386, −0.146]	[−0.325, −0.127]
Infl. Surpr.	13.748***	12.631***
	[12.340, 15.156]	[11.757, 13.505]
Type of infl. surpr.	Point	Mean
F-statistic	221.682	460.780
R^2	0.215	0.231
Observations	101,438	92,765

Note: The table (with panel (a) focusing on 1-year inflation and panel (b) on 3-year inflation) reports results of regressing the interaction of uncertainty and surprise on the interaction of tenure and surprise, as well as inflation surprise, corresponding to the first stage of the 2SLS approach. 95% confidence intervals based on robust standard errors allowing for two-way clustering by respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Reduced Form Regression: Robustness

(a) 1 year inflation

	(1)	(2)	(3)	(4)	(5)	(6)
Dep var:	Δm	Δm	Δm	Δm	Δm	Δm
Tenure × Infl. Surpr.	−0.0061***	−0.0025***	−0.0056***	−0.0027***	−0.0064***	−0.0031***
	[−0.0074, −0.0048]	[−0.0040, −0.0010]	[−0.0068, −0.0044]	[−0.0040, −0.0014]	[−0.0077, −0.0051]	[−0.0046, −0.0016]
Infl. Surpr.	0.1510***	0.0488***	0.1419***	0.0555***	0.1650***	0.0570***
	[0.1334, 0.1686]	[0.0363, 0.0613]	[0.1272, 0.1566]	[0.0430, 0.0680]	[0.1515, 0.1785]	[0.0445, 0.0695]
CPI or PCE	CPI	CPI	PCE	PCE	PCE	PCE
Headline or core	C	C	H	H	C	C
Point or Mean	P	M	P	M	P	M
R^2	0.032	0.003	0.034	0.004	0.036	0.003
Observations	103,888	93,460	103,888	93,460	103,888	93,460

(b) 3 year inflation

	(1)	(2)	(3)	(4)	(5)	(6)
Dep var:	Δm	Δm	Δm	Δm	Δm	Δm
Tenure × Infl. Surpr.	−0.0039***	−0.0021***	−0.0031***	−0.0014**	−0.0040***	−0.0024***
	[−0.0053, −0.0025]	[−0.0036, −0.0006]	[−0.0044, −0.0018]	[−0.0028, −0.0000]	[−0.0055, −0.0025]	[−0.0040, −0.0008]
Infl. Surpr.	0.0887***	0.0383***	0.0800***	0.0356***	0.0964***	0.0441***
	[0.0739, 0.1035]	[0.0224, 0.0542]	[0.0664, 0.0936]	[0.0227, 0.0485]	[0.0829, 0.1099]	[0.0293, 0.0589]
CPI or PCE	CPI	CPI	PCE	PCE	PCE	PCE
Headline or core	C	C	H	H	C	C
Point or Mean	P	M	P	M	P	M
R^2	0.009	0.001	0.010	0.002	0.011	0.002
Observations	101,438	92,765	101,438	92,765	101,438	92,765

Note: Robustness table for the reduced form regressions reported in table 3. Across columns, inflation surprises vary along three dimensions: (i) CPI vs. PCE inflation, (ii) headline vs. core, and (iii) point estimate vs. mean of the belief distribution. 95% confidence intervals based on robust standard errors allowing for clustering by respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Reduced Form Regression: Robustness, controlling for tenure

(a) 1 year inflation

	(1)	(2)	(3)	(4)	(5)	(6)
Dep var:	Δm	Δm	Δm	Δm	Δm	Δm
Tenure × Infl. Surpr.	−0.0053***	−0.0021***	−0.0049***	−0.0023***	−0.0059***	−0.0027***
	[−0.0068, −0.0038]	[−0.0035, −0.0007]	[−0.0063, −0.0035]	[−0.0036, −0.0010]	[−0.0075, −0.0043]	[−0.0042, −0.0012]
Infl. Surpr.	0.1456***	0.0456***	0.1373***	0.0527***	0.1617***	0.0538***
	[0.1271, 0.1641]	[0.0336, 0.0576]	[0.1211, 0.1535]	[0.0396, 0.0658]	[0.1463, 0.1771]	[0.0414, 0.0662]
Tenure	0.0067***	0.0051***	0.0055**	0.0038**	0.0033	0.0035**
	[0.0026, 0.0108]	[0.0021, 0.0081]	[0.0013, 0.0097]	[0.0005, 0.0071]	[−0.0009, 0.0075]	[0.0005, 0.0065]
CPI or PCE	CPI	CPI	PCE	PCE	PCE	PCE
Headline or core	C	C	H	H	C	C
Point or Mean	P	M	P	M	P	M
R^2	0.032	0.003	0.034	0.004	0.036	0.003
Observations	103,888	93,460	103,888	93,460	103,888	93,460

(b) 3 year inflation

	(1)	(2)	(3)	(4)	(5)	(6)
Dep var:	Δm	Δm	Δm	Δm	Δm	Δm
Tenure × Infl. Surpr.	−0.0033***	−0.0020**	−0.0024***	−0.0012	−0.0033***	−0.0023***
	[−0.0049, −0.0017]	[−0.0036, −0.0004]	[−0.0039, −0.0009]	[−0.0026, 0.0002]	[−0.0051, −0.0015]	[−0.0040, −0.0006]
Infl. Surpr.	0.0838***	0.0372***	0.0745***	0.0341***	0.0918***	0.0435***
	[0.0684, 0.0992]	[0.0210, 0.0534]	[0.0599, 0.0891]	[0.0208, 0.0474]	[0.0767, 0.1069]	[0.0281, 0.0589]
Tenure	0.0062***	0.0019	0.0065***	0.0020	0.0046**	0.0007
	[0.0024, 0.0100]	[−0.0015, 0.0053]	[0.0024, 0.0106]	[−0.0015, 0.0055]	[0.0004, 0.0088]	[−0.0026, 0.0040]
CPI or PCE	CPI	CPI	PCE	PCE	PCE	PCE
Headline or core	C	C	H	H	C	C
Point or Mean	P	M	P	M	P	M
R^2	0.009	0.001	0.010	0.002	0.011	0.002
Observations	101,438	92,765	101,438	92,765	101,438	92,765

Note: Same as table A.2, but adding tenure as control. 95% confidence intervals based on robust standard errors allowing for clustering by respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4: Second Stage Regression: Coefficients

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr.	0.0230^{***}	0.0114^{***}	0.0120^{***}	0.0032
	[0.0164, 0.0296]	[0.0049, 0.0179]	[0.0057, 0.0183]	[−0.0019, 0.0083]
Unc. × $\eta_{i,t}$ × Infl. Surpr.	0.0000160^{***}	0.0000075^{**}	0.0000166^{***}	0.0000110^{***}
	[0.0000090, 0.0000230]	[0.0000011, 0.0000139]	[0.0000109, 0.0000223]	[0.0000040, 0.0000180]
Unc. × Infl. Surpr.^{2}	0.0000411	−0.0000025	0.0000212	−0.0000076
	[−0.0000092, 0.0000914]	[−0.0000433, 0.0000383]	[−0.0000173, 0.0000597]	[−0.0000490, 0.0000338]
Infl. Surpr.	-0.0804^{***}	-0.0667^{**}	-0.0458^{**}	−0.0105
	[−0.1242, −0.0366]	[−0.1214, −0.0120]	[−0.0909, −0.0007]	[−0.0552, 0.0342]
Infl. Surpr.^{2}	−0.0017	0.0007	−0.0011	0.0004
	[−0.0045, 0.0011]	[−0.0010, 0.0024]	[−0.0031, 0.0009]	[−0.0009, 0.0017]
$\eta_{i,t}$× Infl. Surpr.	-0.0247^{***}	-0.0128^{***}	-0.0135^{***}	-0.0047^{*}
	[−0.0313, −0.0181]	[−0.0192, −0.0064]	[−0.0199, −0.0071]	[−0.0099, 0.0005]
$\eta_{i,t}$	0.0069^{***}	0.0033^{***}	0.0057^{***}	0.0031^{***}
	[0.0050, 0.0088]	[0.0021, 0.0045]	[0.0043, 0.0071]	[0.0021, 0.0041]
Constant	0.0921^{***}	0.0110	0.0424^{***}	0.0028
	[0.0479, 0.1363]	[−0.0119, 0.0339]	[0.0141, 0.0707]	[−0.0175, 0.0231]
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	103,888	93,460	101,438	92,765

Note: Table reports all the coefficients of the regression in table 4. 95% confidence intervals based on standard errors that account for sampling uncertainty in the first stage and allow for clustering by both respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Second Stage Regression: Robustness

(a) 1-Year Inflation

	Dep. var.: $\Delta m_{i,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Unc. × Infl. Surpr., AME	0.0287***	0.0135***	0.0266***	0.0144***	0.0294***	0.0158***
	[0.0204, 0.0370]	[0.0050, 0.0220]	[0.0193, 0.0339]	[0.0068, 0.0220]	[0.0212, 0.0376]	[0.0072, 0.0244]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0317***	−0.0157***	−0.0289***	−0.0164***	−0.0327***	−0.0184***
	[−0.0400, −0.0234]	[−0.0241, −0.0073]	[−0.0362, −0.0216]	[−0.0239, −0.0089]	[−0.0409, −0.0245]	[−0.0268, −0.0100]
Other controls	Y	Y	Y	Y	Y	Y
CPI or PCE	CPI	CPI	PCE	PCE	PCE	PCE
Headline or core	C	C	H	H	C	C
Point or Mean	P	M	P	M	P	M
Observations	103,888	93,460	103,888	93,460	103,888	93,460

(b) 3-Year Inflation

	Dep. var.: $\Delta m_{i,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Unc. × Infl. Surpr., AME	0.0183***	0.0092**	0.0145***	0.0059*	0.0181***	0.0099***
	[0.0102, 0.0264]	[0.0021, 0.0163]	[0.0073, 0.0217]	[−0.0004, 0.0122]	[0.0099, 0.0263]	[0.0025, 0.0173]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0205***	−0.0115***	−0.0164***	−0.0081**	−0.0205***	−0.0126***
	[−0.0286, −0.0124]	[−0.0187, −0.0043]	[−0.0236, −0.0092]	[−0.0145, −0.0017]	[−0.0287, −0.0123]	[−0.0200, −0.0052]
Other controls	Y	Y	Y	Y	Y	Y
CPI or PCE	CPI	CPI	PCE	PCE	PCE	PCE
Headline or core	C	C	H	H	C	C
Point or Mean	P	M	P	M	P	M
Observations	101,438	92,765	101,438	92,765	101,438	92,765

Note: Same as table 4, but varying the construction of the inflation surprise along three dimensions: (i) CPI vs. PCE inflation, (ii) headline vs. core inflation, (iii) mean of the belief distribution vs. respondents' point estimates. 95% confidence intervals based on standard errors that account for sampling uncertainty in the first stage and allow for clustering by both respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Second Stage Regression: alternative dependent variable

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr., AME	0.0516**	0.0280***	0.0318*	0.0374***
	[0.0068, 0.0964]	[0.0178, 0.0382]	[−0.0003, 0.0639]	[0.0231, 0.0517]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0523**	−0.0308***	−0.0337**	−0.0400***
	[−0.0973, −0.0073]	[−0.0413, −0.0203]	[−0.0658, −0.0016]	[−0.0547, −0.0253]
Other controls	Y	Y	Y	Y
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	80,822	99,984	78,356	97,263

Note: Analogous to table 4, but using the respondents' point estimate of inflation to construct the dependent variable. 95% confidence intervals based on standard errors that account for sampling uncertainty in the first stage and allow for clustering by both respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Second Stage Regression: Additional Controls

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr., AME	0.0157***	0.0064**	0.0058*	−0.0030
	[0.0089, 0.0225]	[0.0000, 0.0128]	[−0.0005, 0.0121]	[−0.0084, 0.0024]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0314***	−0.0154	−0.0148*	−0.0008
	[−0.0468, −0.0160]	[−0.0357, 0.0049]	[−0.0303, 0.0007]	[−0.0182, 0.0166]
Other controls	Y	Y	Y	Y
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	103,888	93,460	101,438	92,765

Note: Same as table 4, but with a larger set of controls: the Kronecker product of (inflation surprise; uncertainty×inflation surprise) with (a constant, the control function component ($\eta_{i,t}$), the inflation surprise, and $\eta_{i,t}^{inv}$), where $\eta_{i,t}^{inv}$ is the inverse of $(\eta_{i,t} - min(\eta_{i,t}) + 10)$. 95% confidence intervals based on standard errors that account for sampling uncertainty in the first stage and allow for clustering by both respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.8: Second Stage Regression: Respondents with All 12 Interviews

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr., AME	0.0233***	0.0176***	0.0093**	0.0038
	[0.0131, 0.0336]	[0.0065, 0.0287]	[0.0009, 0.0178]	[−0.0030, 0.0106]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0247***	−0.0189***	−0.0104**	−0.0049
	[−0.0349, −0.0145]	[−0.0299, −0.0079]	[−0.0190, −0.0019]	[−0.0119, 0.0020]
Other controls	Y	Y	Y	Y
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	63,230	61,268	61,865	60,858

Note: Same as table 4, but restricting the sample to respondents who remain in the survey for all 12 monthly interviews. This addresses the concern that the results reflect selection, e.g. if respondents with high uncertainty are also more likely to drop out of the survey. 95% confidence intervals based on standard errors that account for sampling uncertainty in the first stage and allow for clustering by both respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: Second Stage Regression: High-Sophistication Subsample

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr., AME	0.0458***	0.0079	0.0262***	−0.0018
	[0.0201, 0.0715]	[−0.0190, 0.0348]	[0.0087, 0.0438]	[−0.0160, 0.0124]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0474***	−0.0088	−0.0284***	−0.0000
	[−0.0734, −0.0214]	[−0.0359, 0.0184]	[−0.0462, −0.0106]	[−0.0144, 0.0144]
Other controls	Y	Y	Y	Y
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	28,927	25,148	28,543	25,134

Note: Same as table 4, but restricting the sample to respondents with a college degree or higher, annual household income of $75,000 or more (roughly the median US household income), and a correct answer to the compound-interest numeracy question in the survey. This screens out unsophisticated respondents who report extreme or erratic beliefs: average uncertainty in this subsample is less than half that of the full sample. In all four columns, the confidence intervals include the corresponding full-sample estimates of table 4. 95% confidence intervals based on standard errors that account for sampling uncertainty in the first stage and allow for clustering by both respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.10: Second Stage Regression: 2SLS Estimates (Robustness Across Inflation Measures)

(a) 1-Year Inflation

	Dep. var.: $\Delta m_{i,t}$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Unc. × Infl. Surpr.	0.0143***	0.0116***	0.0214***	0.0160*	0.0199***	0.0155***	0.0282***	0.0206*
	[0.0086, 0.0200]	[0.0034, 0.0198]	[0.0097, 0.0331]	[−0.0011, 0.0331]	[0.0101, 0.0297]	[0.0038, 0.0272]	[0.0088, 0.0476]	[−0.0008, 0.0420]
Infl. Surpr.	−0.0866***	−0.0975**	−0.1874**	−0.1661	−0.1607**	−0.1409**	−0.3016**	−0.2212
	[−0.1488, −0.0244]	[−0.1898, −0.0052]	[−0.3485, −0.0263]	[−0.3831, 0.0509]	[−0.2903, −0.0311]	[−0.2797, −0.0021]	[−0.5942, −0.0090]	[−0.4956, 0.0532]
CPI or PCE	CPI	CPI	CPI	CPI	PCE	PCE	PCE	PCE
Headline or core	H	H	C	C	H	H	C	C
Point or Mean	P	M	P	M	P	M	P	M
Observations	103,888	93,460	103,888	93,460	103,888	93,460	103,888	93,460

(b) 3-Year Inflation

	Dep. var.: $\Delta m_{i,t}$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Unc. × Infl. Surpr.	0.0094***	0.0032	0.0165**	0.0084**	0.0142**	0.0065	0.0216*	0.0105**
	[0.0028, 0.0160]	[−0.0023, 0.0087]	[0.0029, 0.0301]	[0.0003, 0.0165]	[0.0023, 0.0261]	[−0.0013, 0.0143]	[−0.0009, 0.0441]	[0.0010, 0.0200]
Infl. Surpr.	−0.0687*	−0.0166	−0.1766*	−0.0785	−0.1363	−0.0494	−0.2678	−0.1019*
	[−0.1461, 0.0087]	[−0.0769, 0.0437]	[−0.3722, 0.0190]	[−0.1747, 0.0177]	[−0.3002, 0.0276]	[−0.1382, 0.0394]	[−0.6225, 0.0869]	[−0.2180, 0.0142]
CPI or PCE	CPI	CPI	CPI	CPI	PCE	PCE	PCE	PCE
Headline or core	H	H	C	C	H	H	C	C
Point or Mean	P	M	P	M	P	M	P	M
Observations	101,438	92,765	101,438	92,765	101,438	92,765	101,438	92,765

Note: This table reports two-stage least squares (2SLS) estimates of the second-stage regression of $\Delta m_{i,t}$ (measured based on the bin mean) on the interaction of uncertainty and inflation surprises, Unc. × $\widehat{\pi}_{i,t}$, and on the inflation surprise $\widehat{\pi}_{i,t}$. The endogenous regressor Unc. × $\widehat{\pi}_{i,t}$ is instrumented using tenure interacted with the inflation surprise, i.e., Tenure$_{i,t}$ × $\widehat{\pi}_{i,t}$, with the inflation surprise included as an exogenous regressor. Panel (a) uses 1-year expectations, panel (b) uses 3-year expectations. Columns vary the construction of the inflation surprise along three dimensions: (i) CPI vs. PCE inflation, (ii) headline vs. core inflation, and (iii) point forecast vs. lagged bin mean. 95% confidence intervals based on standard errors two-way clustered by respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.11: First stage regressions with 1/tenure

(a) First Stage, 1 year inflation

	(1)	(2)	(3)	(4)
Dep var:	Unc.	Unc.	Unc.	Unc.
1/Tenure	10.553***	9.794***	5.433***	7.420***
	[9.398, 11.707]	[8.016, 11.573]	[4.419, 6.447]	[5.630, 9.210]
Infl. Surpr.			−0.601***	0.178***
			[−0.757, −0.445]	[0.078, 0.278]
Type of infl. surpr.			Point	Mean
F-statistic	326.557	118.552	96.839	39.238
R^2	0.005	0.002	0.030	0.003
Observations	116,720	104,595	103,888	93,460

(b) First Stage, 3 year inflation

	(1)	(2)	(3)	(4)
Dep var:	Unc.	Unc.	Unc.	Unc.
1/Tenure	10.638***	10.764***	5.102***	8.892***
	[9.455, 11.820]	[8.943, 12.585]	[3.980, 6.223]	[7.130, 10.655]
Infl. Surpr.			−0.600***	0.272***
			[−0.755, −0.445]	[0.177, 0.368]
Type of infl. surpr.			Point	Mean
F-statistic	316.331	136.546	91.272	62.197
R^2	0.005	0.002	0.027	0.005
Observations	117,019	104,827	101,438	92,765

Note: Same as table 2, but with the instrument entering as 1/tenure instead of linearly. This is the functional form suggested by Bayesian learning, under which the posterior variance declines with the inverse of the number of signals received; figure A.3 shows that it closely tracks the average uncertainty by tenure bin. 95% confidence intervals based on robust standard errors allowing for two-way clustering by person and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.12: Second Stage Regression: 1/tenure Instrument

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr., AME	0.0219^{***}	0.0118^{***}	0.0130^{***}	0.0029
	[0.0154, 0.0284]	[0.0050, 0.0186]	[0.0066, 0.0193]	[−0.0026, 0.0084]
$\eta_{i,t}$ × Infl. Surpr., AME	-0.0235^{***}	-0.0131^{***}	-0.0144^{***}	−0.0043
	[−0.0301, −0.0169]	[−0.0198, −0.0063]	[−0.0208, −0.0079]	[−0.0099, 0.0013]
Other controls	Y	Y	Y	Y
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	103,888	93,460	101,438	92,765

Note: Same as table 4, but the control-function first stage regresses uncertainty on 1/tenure in place of tenure (see table A.11 and figure A.3). 95% confidence intervals based on standard errors that account for sampling uncertainty in the first stage and allow for clustering by both respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.13: Second Stage Regression: Previous Month's Realized Inflation as Surprise

(a) CPI

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr., AME	0.0228***	0.0102***	0.0128***	0.0074**
	[0.0152, 0.0304]	[0.0039, 0.0166]	[0.0057, 0.0200]	[0.0014, 0.0134]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0248***	−0.0119***	−0.0145***	−0.0091***
	[−0.0324, −0.0171]	[−0.0183, −0.0054]	[−0.0217, −0.0073]	[−0.0151, −0.0031]
Other controls	Y	Y	Y	Y
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	103,888	93,460	101,438	92,765

(b) PCE

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
Unc. × Infl. Surpr., AME	0.0262***	0.0127***	0.0153***	0.0091***
	[0.0184, 0.0340]	[0.0051, 0.0202]	[0.0076, 0.0229]	[0.0024, 0.0158]
$\eta_{i,t}$ × Infl. Surpr., AME	−0.0289***	−0.0149***	−0.0174***	−0.0115***
	[−0.0367, −0.0211]	[−0.0224, −0.0074]	[−0.0250, −0.0097]	[−0.0182, −0.0048]
Other controls	Y	Y	Y	Y
Type of infl. surpr.	Point	Mean	Point	Mean
Observations	103,888	93,460	101,438	92,765

Note: Same as table 4, but the inflation surprise is constructed from the *previous* month's realized inflation—the number actually released during the survey month—rather than the current month's, which is not yet available at the time of the survey. The lags of expectations are unchanged: the surprise is $12\pi_{t-1}$ minus the point forecast (lagged one month, columns (1) and (3)) or the bin mean (lagged two months, columns (2) and (4)). Panel (a) constructs the surprise from headline CPI inflation and panel (b) from headline PCE inflation. 95% confidence intervals based on standard errors that account for sampling uncertainty in the first stage and allow for clustering by both respondent and time are reported in brackets. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.14: Second Stage Regression: Alternative Surprise Definitions

	1 year inflation		3 year inflation	
	(1)	(2)	(3)	(4)
1. Baseline: $\pi_t - m_{i,t-j}$	0.0229^{***}	0.0114^{***}	0.0120^{***}	0.0032
	[0.0163, 0.0295]	[0.0049, 0.0179]	[0.0057, 0.0183]	[−0.0019, 0.0083]
2. Consensus mean: $\pi_t - \bar{m}_{t-j}$	0.0157^{***}	0.0069^{*}	0.0079^{***}	-0.0043^{**}
	[0.0096, 0.0218]	[−0.0000, 0.0139]	[0.0022, 0.0137]	[−0.0080, −0.0006]
3. Consensus median: $\pi_t - \text{med}(m)_{t-j}$	0.0123^{***}	0.0059	0.0051^{*}	-0.0051^{***}
	[0.0054, 0.0192]	[−0.0015, 0.0134]	[−0.0009, 0.0112]	[−0.0090, −0.0012]
4. Full-sample mean: $\pi_t - \bar{m}$	0.0107^{***}	0.0048	0.0080^{***}	-0.0032^{*}
	[0.0048, 0.0165]	[−0.0018, 0.0114]	[0.0031, 0.0129]	[−0.0064, 0.0000]
5. Full-sample median: $\pi_t - \text{med}(m)$	0.0034	0.0026	0.0031	-0.0044^{**}
	[−0.0038, 0.0105]	[−0.0047, 0.0100]	[−0.0023, 0.0086]	[−0.0079, −0.0008]
6. Tenure residualized: $\pi_t - \tilde{m}_{i,t-j}$	0.0122^{***}	0.0054^{***}	0.0036^{*}	-0.0007
	[0.0077, 0.0166]	[0.0013, 0.0095]	[−0.0004, 0.0077]	[−0.0040, 0.0026]
Type of infl. surpr.	Point	Mean	Point	Mean
Observations (rows 1–5)	103,888	93,460	101,438	92,765
Observations (row 6)	103,875	93,449	101,418	92,744

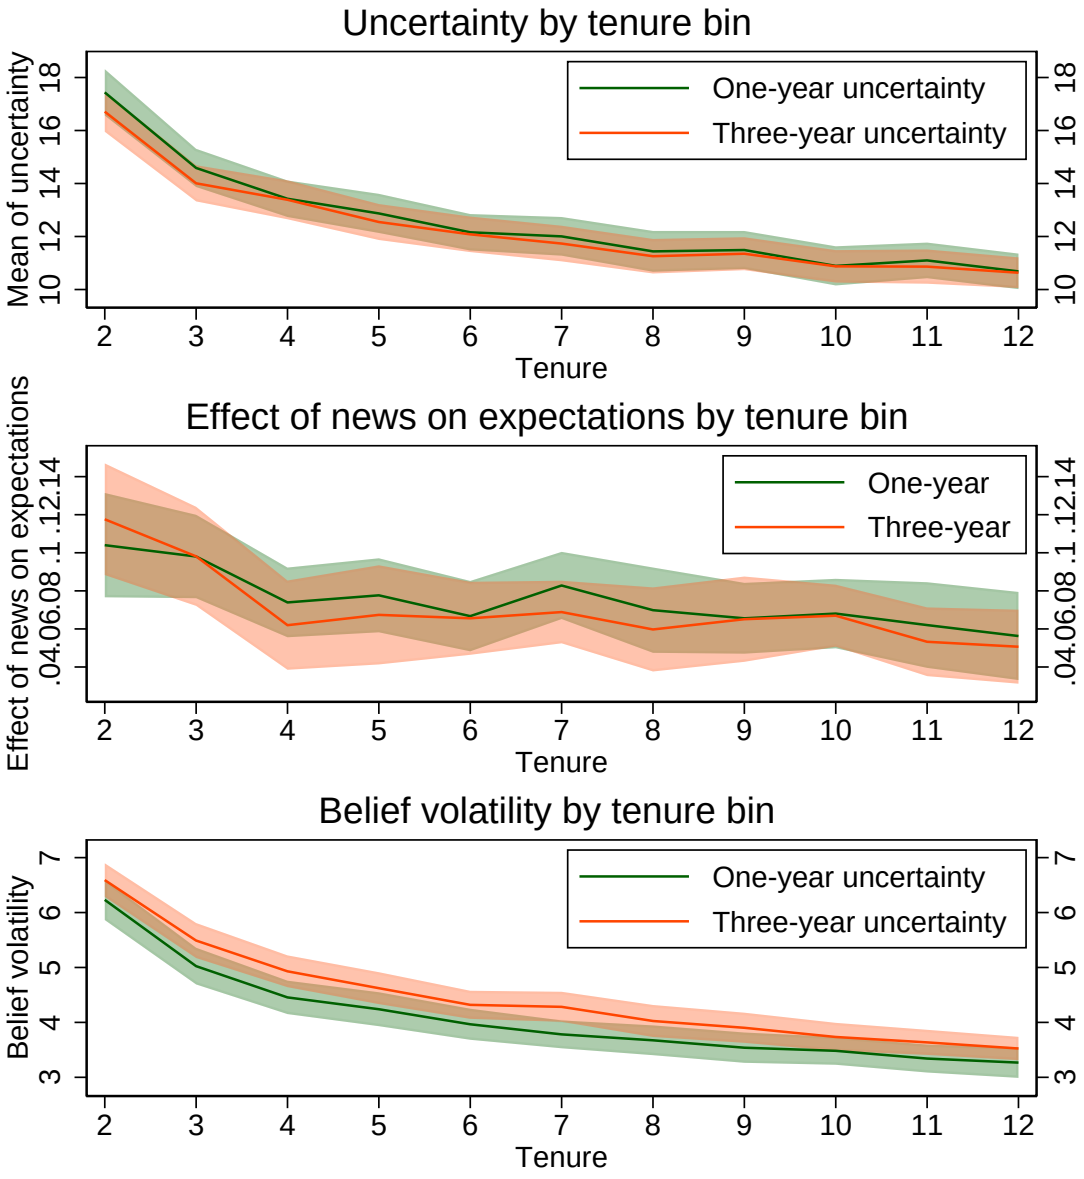
Note: Each row reports the average marginal effect (AME) of uncertainty × inflation surprise – the estimate of $E[\sigma^{-2}]$ – from the control-function second stage of table 4, with the agent's own lagged expectation in the surprise replaced by an aggregate counterpart. Row 1 repeats the baseline surprise $\hat{\pi}_{i,t} = \pi_t - m_{i,t-j}$ ($j = 1$ for the point forecast, $j = 2$ for the bin mean). Rows 2–3 replace $m_{i,t-j}$ with the lagged cross-sectional mean ($\bar{m}_{t-j}$) or median ($\text{med}(m)_{t-j}$) of the same expectation measure, with the same lag conventions. Rows 4–5 replace it with the full-sample mean ($\bar{m}$) or median ($\text{med}(m)$) of the expectation measure – the same constant for all agents and periods, so that the surprise varies only through realized inflation. Row 6 residualizes expectations with respect to tenure dummies. Rows 1–5 are estimated on the baseline samples of table 4; the alternative surprises are defined only where the baseline surprise is available. Columns (1)–(2) use the 1-year expectations and columns (3)–(4) the 3-year expectations, with the surprise measured based on the point forecast (Point) or the bin mean (Mean). 95% confidence intervals in brackets account for sampling uncertainty in the first stage and allow for clustering by both respondent and time. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.15: Second Stage Regression for $(\Delta m)^2$: Measurement Adjustments

	1-year inflation		3-year inflation	
	(1)	(2)	(3)	(4)
1. Baseline: $(\Delta m_t)^2$	0.1336***	0.0763***	0.1902***	0.0753***
	[0.1038, 0.1634]	[0.0661, 0.0864]	[0.1493, 0.2311]	[0.0655, 0.0852]
2. Tenure-demeaned: $\left(\Delta m_t - \overline{\Delta m}_\tau\right)^2$	0.1333***	0.0773***	0.1882***	0.0767***
	[0.1039, 0.1627]	[0.0671, 0.0876]	[0.1478, 0.2285]	[0.0667, 0.0866]
3. Noise-corrected: $\Delta m_t \sum_{s=-1}^{1} \Delta m_{t+s}$	0.1550	0.0556***	0.1077*	0.0380***
	[−0.1299, 0.4398]	[0.0250, 0.0862]	[−0.0047, 0.2201]	[0.0184, 0.0576]
4. Tenure-demeaned and noise-corrected	0.1526	0.0571***	0.1028*	0.0399***
	[−0.1237, 0.4289]	[0.0266, 0.0876]	[−0.0032, 0.2089]	[0.0202, 0.0596]
Type of inflation expectation	Point	Mean	Point	Mean
Observations (rows 1–2)	102,140	116,720	101,059	117,019
Observations (rows 3–4)	63,995	74,550	63,368	74,803

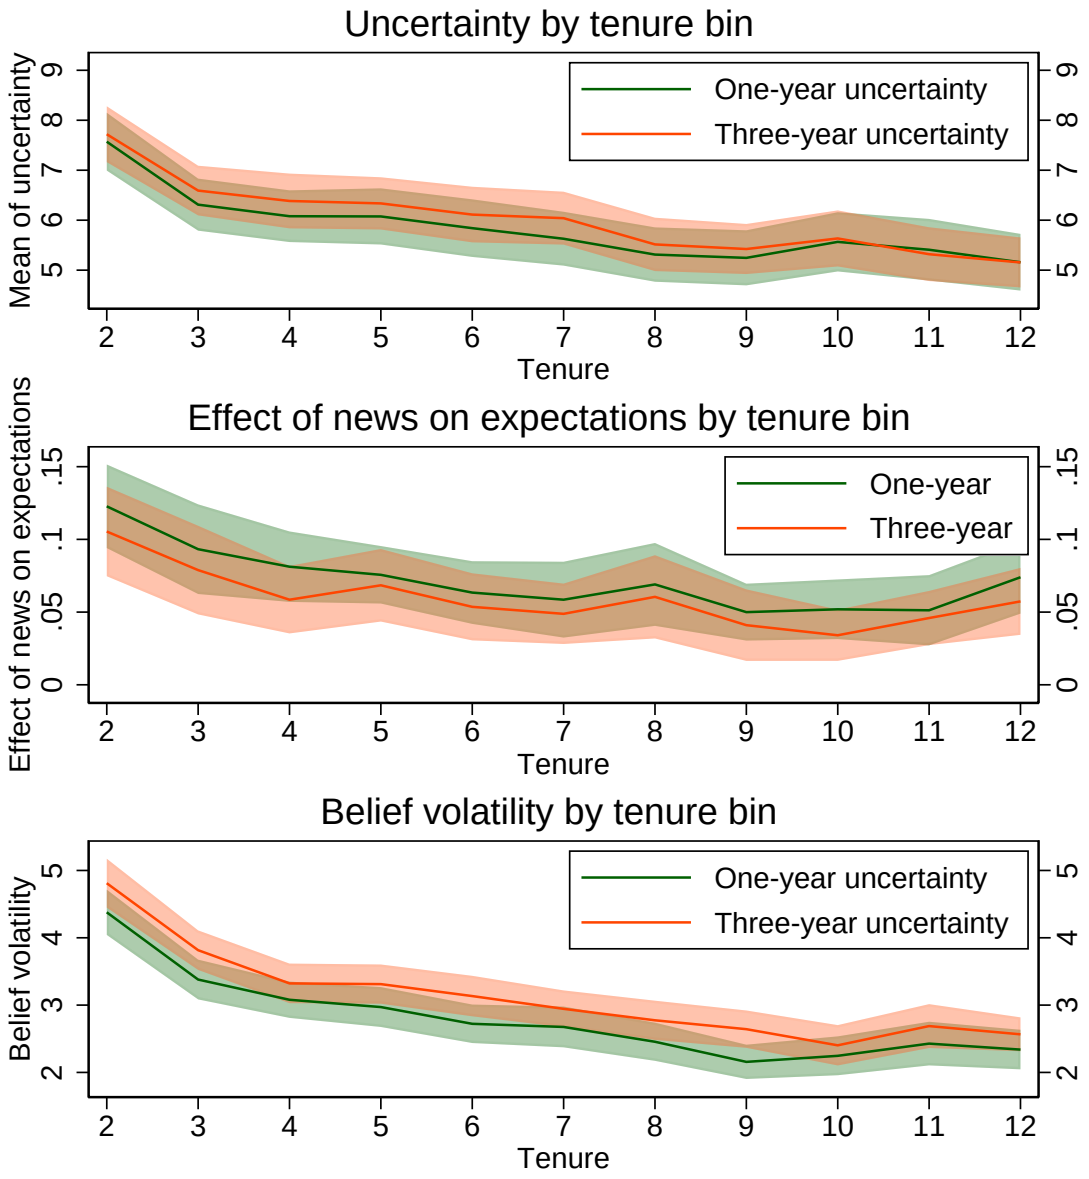
Note: Each row reports the AME (average marginal effect) of $Unc_{i,t}^2$ from the second-stage regression of table 5, with the dependent variable $(\Delta m_{i,t})^2$ replaced by the variant in the row label; uncertainty, the first stage, and all control-function terms are unchanged. Row 1 repeats the baseline. Row 2 demeans the change in expectations by its tenure-bin mean ($\overline{\Delta m}_\tau$) before squaring. Row 3 applies the noise correction described in appendix A.2, replacing $(\Delta m_{i,t})^2$ with $\Delta m_{i,t}(\Delta m_{i,t-1} + \Delta m_{i,t} + \Delta m_{i,t+1})$, whose expectation purges serially uncorrelated reporting noise; it requires three consecutive monthly revisions, which reduces the sample. Row 4 combines the two adjustments. Columns (1)–(2) use 1-year inflation expectations and columns (3)–(4) use 3-year inflation expectations; within each horizon, the dependent variable is constructed using point expectations (Point) or bin means (Mean). 95% confidence intervals in brackets account for sampling uncertainty in the first stage and allow for clustering by both respondent and time. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A.1: First stage and reduced-form relationships: respondents with all 12 interviews



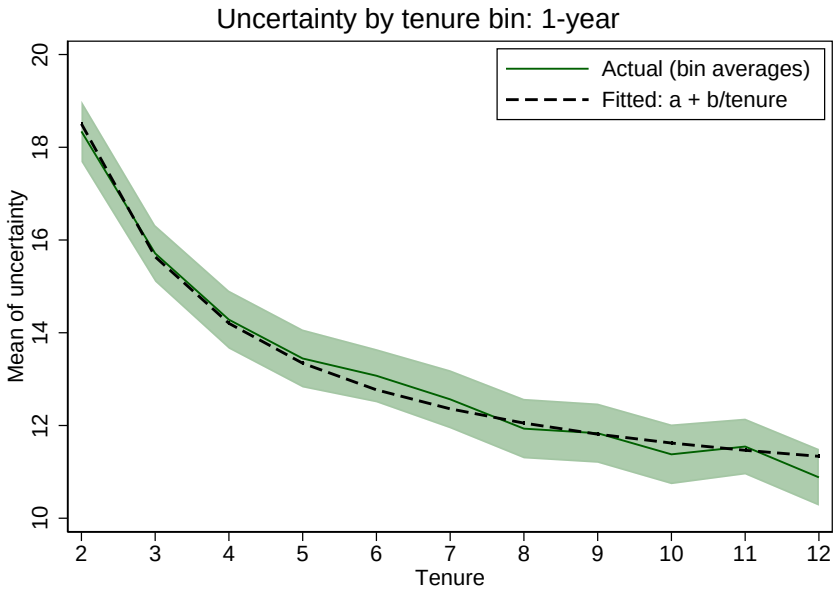
Note: Same as figure 3, but restricting the sample to respondents who remain in the survey for all 12 monthly interviews. In each panel, the lines are point estimates and the shaded regions 95-percent confidence bands. Standard errors are two-way clustered by person and time. The top panel reports coefficients from a regression of uncertainty (each agent's variance) on tenure indicators. The middle panel is from a regression of the change in each agent's conditional mean on inflation news interacted with tenure indicators. The bottom panel is from a regression of the squared change in expectations on tenure indicators.

Figure A.2: First stage and reduced-form relationships: high-sophistication subsample

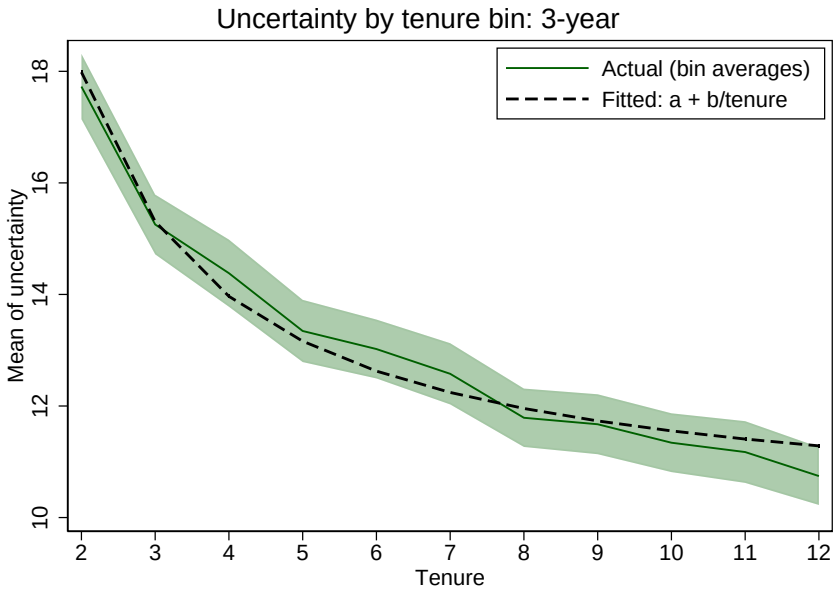


Note: Same as figure 3, but restricting the sample to respondents with a college degree or higher, annual household income of $75,000 or more (roughly the median US household income), and a correct answer to the compound-interest numeracy question in the survey (see table A.9). In each panel, the lines are point estimates and the shaded regions 95-percent confidence bands. Standard errors are two-way clustered by person and time. The top panel reports coefficients from a regression of uncertainty (each agent's variance) on tenure indicators. The middle panel is from a regression of the change in each agent's conditional mean on inflation news interacted with tenure indicators. The bottom panel is from a regression of the squared change in expectations on tenure indicators.

Figure A.3: Uncertainty by tenure: actual bin averages vs. fitted values from 1/tenure



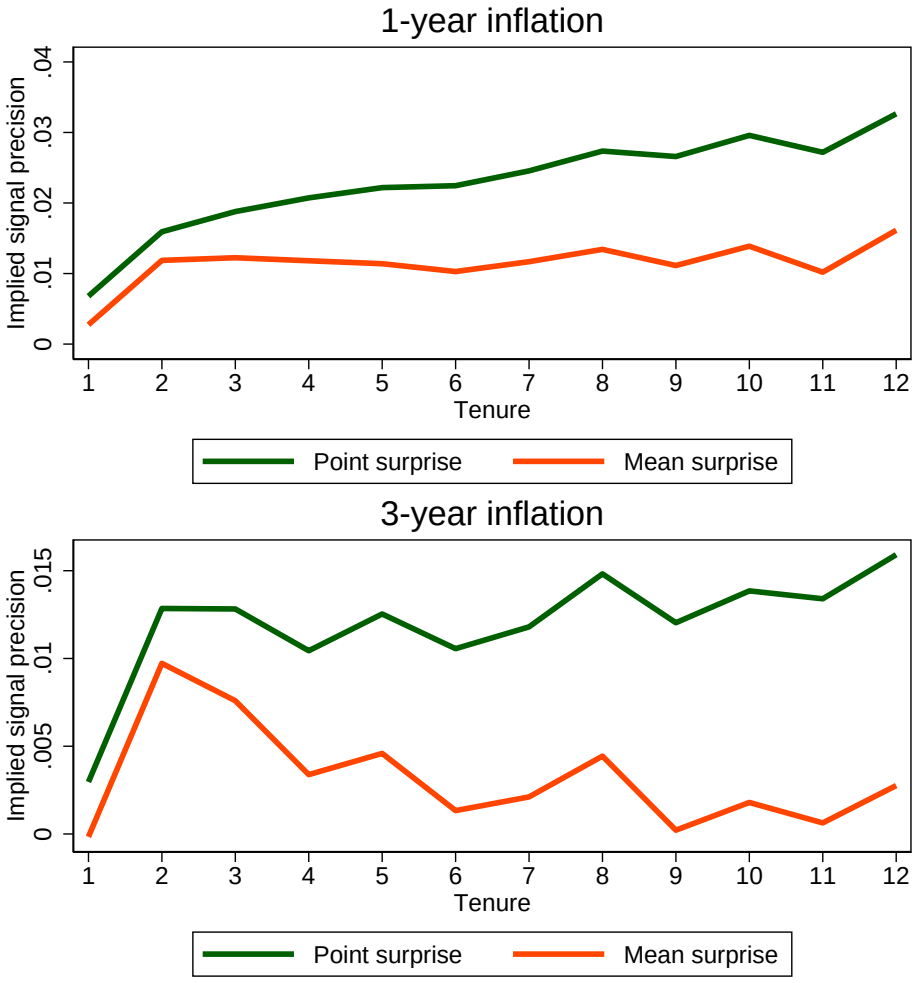
(a) 1-year inflation



(b) 3-year inflation

Note: The solid lines reproduce the top panel of figure 3: average uncertainty by tenure bin (shaded regions are 95-percent confidence bands, two-way clustered by person and time), separately for 1-year (panel a) and 3-year (panel b) expectations. The dashed lines are the fitted values from a regression of uncertainty on 1/tenure on the same sample, the first stage used in tables A.11 and A.12.

Figure A.4: Implied signal precision by months of tenure



Note: Signal precision σ_n^{-2} implied by equation (24) at each month of tenure n, computed from average uncertainty by tenure bin in the full sample (see appendix A.1). For each horizon and each inflation-surprise measure (point forecast or bin mean), the steady-state precision σ_1^{-2} is calibrated so that the average of σ_n^{-2} over the twelve tenure months equals the corresponding estimate of $\mathbb{E}\left[\sigma^{-2}\right]$ (the Unc. $\times$ Infl. Surpr. AME) in table 4. The top panel refers to 1-year inflation expectations and the bottom panel to 3-year expectations.