

Using density forecasts to measure the stability of inflation expectations*

Ian Dew-Becker, Stefano Giglio, and Pooya Molavi

March 5, 2026

Abstract

The anchoring of inflation expectations is a major objective of modern monetary policy. This paper estimates variation over time in the sensitivity of inflation expectations to news. The function describing the response of an agent's expectations to news is fundamentally unobservable since on a given date we only see the single realization of news that actually occurred. However, under the assumption that agents apply Bayes' rule and that they believe their signals are Gaussian (hence allowing for a wide range of behavioral biases), the marginal response of expectations to signals is proportional to agents' uncertainty about future inflation. Empirically, both in the time-series and the cross-section, reported uncertainty both contemporaneously explains and also predicts the future sensitivity of expectations to news. The results imply that as of 2025, inflation expectations are 2–3 times more sensitive to news than prior to 2020.

One of the primary goals of modern monetary policy is ensuring that inflation expectations remain well anchored. Anchoring is taken in some loose sense to mean that expectations are not too sensitive to news and therefore do not fluctuate very much. One can think of backward-looking, contemporaneous, and forward-looking measures of anchoring. A backward-looking measure would simply ask whether inflation expectations were insensitive to news and stable over some period. Contemporaneously, the question is how sensitive expectations are to news today: what is the slope of the function mapping today's signals to expected future inflation? A forward-looking measure then would try to forecast that slope in the future.

*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. The paper benefited greatly from discussions with Jeff Campbell.

The core difficulty 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 expectations to news only having observed the actual news that was realized. Just measuring the past volatility of changes in expectations is also not enough. First, a backward-looking measure like that does not directly answer the question of how stable expectations are going forward. More importantly, though, expectations could have been volatile either because they were sensitive to news, or because the news was extreme.¹ This paper’s goal is to ask how well anchored expectations are by measuring, *in real time*, how sensitive expectations are to pieces of information – such as inflation announcements – holding their content (i.e. the size of the surprise) fixed.

Contribution

Measuring how inflation expectations will react to future news necessarily requires thinking about how people would change their expectations depending on what actually happens in the future. That is, it involves estimating counterfactuals, which fundamentally requires imposing structure (or, as in [Armantier et al. \(2022\)](#), discussed further below, one can directly ask people about hypotheticals). This paper imposes a narrow structural assumption, which is simply that agents update the expectations using Bayes’ formula and they assume their signals about future inflation have normally distributed errors. Note, critically, the assumption *is not* that they are fully rational Bayesians – we discuss a wide range of behavioral biases that are allowed within the analysis.

Given the setup, the paper’s core theoretical contribution is to show that the mapping from signals to beliefs can be easily and intuitively recovered from an agent’s posterior distribution. A particularly nice feature of that mapping is that the derivative of expectations with respect to signals is proportional to agents’ posterior variance – which we refer to as uncertainty – over future inflation.

The paper then empirically evaluates how well that implication actually describes the dynamics of beliefs in a panel survey. It provides evidence from instrumental variables methods supporting the model’s core implication that periods when agents report relatively high uncertainty should be periods in which their expectations are more sensitive to signals and more volatile. In other words, agents’ reported uncertainty measures how well anchored their expectations are.

Finally, we examine data on uncertainty over time in the US, focusing on the post-2013 period, and find that as of 2025, even after falling significantly since its peak in 2022,

¹See [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. \(2022\)](#)

uncertainty remained 2–3 times higher than its values prior to 2020, implying expectations are 2–3 more sensitive to a given shock, and hence 1/3–1/2 as well anchored.

Methods and results

The paper’s basic approach is to think about 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 – can be expressed as a power series in terms of the realized signal scaled by its precision. [Dytso, Poor and Shitz \(2022\)](#) show that the coefficients in the power series are actually equal to the agents’ posterior cumulants. The first derivative is the posterior variance, the second the posterior third moment, and the n th is the posterior $(n + 1)$ th cumulant.

The fundamental goal in quantifying anchoring is to estimate an unobservable function: the mapping from signals to posterior expectations. The filtering theorem shows that there is a duality: the posterior distribution can be transformed to yield the mapping from signals to expectations.

All of that depends on the assumption that agents apply some form of Bayes’ formula under the belief that their signals are normally distributed. That assumption is certainly not literally true, and it might not even be approximately true, so the paper’s next task is to evaluate how well the model’s duality prediction actually describes the data. The fact that the response function is not observable means that it is not possible to directly test the model (or any model of the response function, for that matter). The first-order prediction, though, is that agents’ beliefs are more sensitive to signals when their uncertainty is higher, and we can evaluate that in the data in the sense of asking whether their beliefs in those periods covary more strongly with inflation surprises and are more volatile.

While there are a number of data sources with information on inflation expectations, the paper focuses primarily on 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 distributions for future inflation outcomes. The main drawback of the SCE, on the other hand, is that the inflation forecasts for which there is a long time series are only at the one- and three-year horizons, whereas policymakers often focus more on longer-term expectations.

From a time-series perspective, periods in which people in the SCE on average report

higher uncertainty are also periods in which reported average expectations covary more strongly with surprises in realized inflation and are more volatile. However, there is an identification problem, which is that the precision of signals – which is not observable – likely varies over time and across people and is inevitably correlated with their uncertainty. Low signal precision mechanically drives future uncertainty up. In the other direction, high uncertainty might cause agents to gather more information, increasing signal precision. So the two variables influence each other, and even ignoring structural relationships, both variables could obviously be simultaneously driven by some third factor.

We therefore use the panel dimension of the data to run an instrumental variables analysis. We argue, based on the evidence of [Kim and Binder \(2023\)](#), that tenure in the survey is a useful instrument for uncertainty. Over the course of respondents' time in the survey, their reported uncertainty about future inflation declines, showing that the instrument is relevant. 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 exogeneity assumption is that attention is orthogonal to 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 fact that uncertainty declines monotonically suggests that attention remains at least somewhat stable across months of tenure. 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 \(2011\)](#).

Everything so far has been about contemporaneous anchoring – understanding the model's prediction that the sensitivity of expectations to signals today depends on today's posterior cumulants. For many practical purposes, though, what matters more is the sensitivity of expectations to signals *going forward*. That is much more of a forecasting problem than a

structural identification problem. The question is simply what variables can help us predict the future sensitivity of expectations to signals. The model implies that the way to do that is by forecasting uncertainty and signal precision.

Empirically, the analysis finds that current uncertainty does in fact have forecasting power for the future behavior of expectations. As a specific example, SCE inflation expectations became much more volatile and sensitive to inflation news in 2021 and 2022, exactly when uncertainty was high. Interestingly, uncertainty rose *before* expectations did. Furthermore, while uncertainty reported in the SCE has come back down from its peak, even as of December, 2025 it was at last twice as high as its value prior to 2020.

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 Another complementary approach to what is in this paper is in [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). Our analysis implies that the high-frequency responsiveness that they measure could be further conditioned on uncertainty (or tenure in the survey).

[Carvalho et al. \(2023\)](#) is also very closely related. That paper 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, whereas [Carvalho et al. \(2023\)](#) assumes they are boundedly rational. That paper also estimates a fully specified model for inflation and expectations, whereas the nature of the more general setup here means that it does not yield a full description of dynamics.

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,](#)

²See 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.

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, and studying that would be a straightforward extension of the analysis.

More generally, Reis (2025) reviews recent evidence on inflation expectations and their link to post-Covid inflation. This paper builds closely on that work by asking what factors drive expectations. Reis (2025) emphasizes the causal effects that expectations can have. The evidence this paper presents that expectations are now less firmly anchored than prior to Covid then suggests that inflation itself may in the future be more volatile. 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 4 studies the model’s implications for forecasting expectational anchoring going forward and section 5 concludes.

1 Theory

The aim of this section is to understand how to measure the response of an agent’s beliefs about inflation to signals they observe. Specifically, given a signal y_t , with π_t^* being the object agents are learning about, we want to understand properties of the function

$$E[\pi_t^* | y_t] \tag{1}$$

where E is a *subjective* expectation operator. There is no assumption yet about its rationality or any of its other properties.

1.1 Economic environment

The log price level is denoted by p_t . H -period inflation starting at date t is

$$\pi_{t,t+H} = p_{t+H} - p_t \tag{2}$$

The survey data that we use asks agents about inflation over 12-month periods starting immediately and 24 months in the future, so that $\pi_t^* = \pi_{t+j,t+j+H}$, with $j \in \{0, 24\}$.

Assumption 1 π_t^* follows an arbitrary (discrete-time) stochastic process. In particular, is not necessarily linear, Gaussian, or homoskedastic.

Given the observed dynamics of inflation, assuming π_t^* is normal here would fail to fit many features of the data.

We assume that agents receive signals about future inflation.

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

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

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

In practice the signal is agent-specific, and below we add an i subscript to indicate different agents.

To be clear, agents are modeled here as receiving information about the future. y_t gives them information about inflation that will be realized in future months. An infinitely precise signal would represent a crystal ball or oracle that perfectly reveals the future.

The assumption that agents treat the signals as conditionally normally distributed is somewhat restrictive (if very standard), but necessary for the main theorem we use below.^{4,5} Note, though, that we neither assume that the signal is *truly* normally distributed nor that its mean and variance are actually π_t^* and σ_t^2 .

1.2 A duality result for updating

The inflation forecasting problem here is a filtering problem: π_t^* is a dynamic latent variable about which agents receive signals. The usual Kalman filter does not apply, though, because we allow for arbitrary dynamics for π_t^* . The two general filtering steps are prediction – going from a date- $(t - 1)$ posterior to a date- t prior – and an update – going from the date- t prior to the date- t posterior (conditional on the date- t signal) using Bayes’ formula.

Assumption 3 *On all dates agents have 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).*

⁴One way to motivate it 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.

⁵There are other versions of theorem 1 below for alternative error distributions for the signal. In general a result can be derived when the distribution is in an exponential family. We use the normal distribution here just because it is a standard benchmark and natural if agents observe many independent sources of information.

Assumption 3 says that agents are rational to the extent that their beliefs can be described by probability distributions that they update via Bayes' formula. It does not impose though, that their specification for either the signals (assumption 2) or for the dynamics of π_t^* is correct. Appendix B.1 describes a range of deviations from full rationality that the analysis allows.

Definition 1 $E[x | y^t]$ denotes the expectation and $\kappa_j(x | y^t)$ denotes the j th cumulant of x given the history of signals observed up to date t , $y^t \equiv \{y_t, y_{t-1}, \dots\}$, **under the agent's subjective probability distribution.**

Recall that the first three cumulants are identically equal to the first three central moments. While the following result uses all the cumulants, we ultimately only discuss the first three.

Theorem 1 [Dytso, Poor and Shitz (2022) theorem 4] Under assumptions 2 and 3, there exists a neighborhood of any point \bar{Y} such that for all Y in that neighborhood,

$$E[\pi_t^* | y_t = Y, y^{t-1}] = \sum_{j=0}^{\infty} \frac{\kappa_{j+1}(\pi_t^* | y_t = \bar{Y}, y^{t-1})}{j!} \left(\frac{Y - \bar{Y}}{\sigma_t^2} \right)^j \quad (4)$$

See appendix A for proofs.

One way to think about equation (4) is to take \bar{Y} to be the value of the signal that was actually realized and Y to be a counterfactual value. It then shows that a feasible way to measure our fundamental object of interest, the posterior mean for *counterfactual* values of the signal, is with a power series in which the coefficients are the realized posterior cumulants. That is the key result driving the paper's analysis and what we will aim to test empirically.

To simplify the notation going forward, we denote the posteriors on date t by

$$E_t[\cdot] \equiv E[\cdot | y^t] \quad (5)$$

$$\kappa_{j,t} \equiv \kappa_j(\pi_t^* | y^t) \quad (6)$$


Corollary 1 *The local sensitivity of expectations to signals is*

$$\frac{d}{dy_t} \kappa_{1,t} = \kappa_{2,t} \sigma_t^{-2} \quad (7)$$

Figure 1 displays an example of the posterior expectation as a function of the signal. The mapping is in general nonlinear, and its slope at each point is equal to the posterior second

moment for that value of the signal. The full behavior (when the power series converges globally) is described by the power series in (4). For relatively low signals, the posterior variance is small and the posterior expectation is relatively insensitive to signals. For high signals, the posterior variance is high and the expectation is much more sensitive.

Figure 1: Hypothetical Bayesian update



Note: The figure assumes a prior distribution $\pi^* \sim \chi^2(3)$ and that $y \sim N(\pi^*, (1/2)^2)$. The lines plot the posterior mean and variance for π^* conditional on the observed signal.

1.3 Higher-order updating

Beyond the first moment, the same sort of analysis that proves theorem 1 also yields the following more general result:

Theorem 2 *Under the same assumptions as theorem 1,*

$$\frac{d}{d(y_t \sigma_t^{-2})} \kappa_{j,t} = \kappa_{j+1,t} \quad (8)$$

Specifically, for $j = 2$,

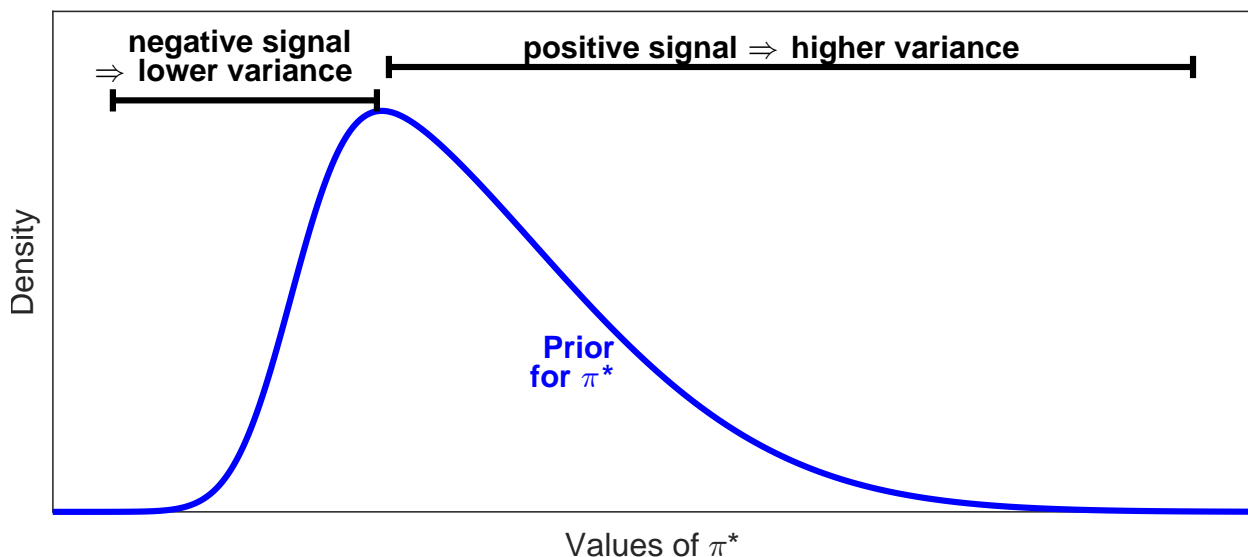
$$\frac{d}{dy_t} \kappa_{2,t} = \kappa_{3,t} \sigma_t^{-2} \quad (9)$$

This result is important in analyzing endogeneity below because it says that $\kappa_{2,t}$ is driven by the same signals that drive expectations, $\kappa_{1,t}$ – uncertainty cannot be taken to be exogenous.

Equation 9 shows that when agents have positively skewed beliefs, so that the right tail of their distribution is longer than the left, positive signals about π_t^* make that long right tail the more likely outcome and raise uncertainty. That is important in thinking about anchoring. If $\kappa_{2,t}$ measures how well anchored beliefs are, then $\kappa_{3,t}$ determines the stability of the anchor.

Figure 2 helps visualize that idea. If an agent has a positively skewed prior, then a positive signal about inflation raises the probability associated with the right-hand side of the distribution, which is more spread out, raising the posterior variance. While the intuition is fairly simple, part of what is surprising about equation (9) is that it really is the third moment alone that determines how the second moment responds to signals.

Figure 2: The effect of updating on the posterior variance



Note: The blue line represents some arbitrary prior distribution for π^* . When signals are relatively high (roughly, above the prior mean), the right-hand side of the distribution becomes more likely and thus agents' conditional variance for π^* rises. When signals are low, the conditional variance falls.

2 Empirically evaluating the relationship between uncertainty and sensitivity

This section develops an approach to empirically testing the predictions from section 1. The key question is whether, all else equal, higher uncertainty is associated with greater sensitivity of expectations to signals. The analysis shows that in general

regressions of the sensitivity of expectations to news on uncertainty are biased and then shows how that bias can be eliminated with an instrument for uncertainty.

2.1 Information structure using inflation as a signal

In taking the model to the data, we want to ask how beliefs respond to realized inflation, but agents obviously have many sources of information beyond just aggregate inflation. This section therefore describes a slight generalization of the analysis above that allows agents to observe two signals – realized inflation and a second factor combining all their unobserved sources of δ . In order to rationalize treating realized inflation, π_t , as a signal, we assume just for this part of the analysis the following:

Assumption 4 *Inflation follows a trend-plus-noise model of the form*

$$\pi_t \sim N(\pi_t^*, \sigma_{\pi,t}^2) \quad (10)$$

Assumption 5 *Agents receive two signals. The first, $s_{i,t}$, is related to realized inflation and is distributed as*

$$s_{i,t} \sim N(\pi_t, \sigma_{s,i,t}^2) \quad (11)$$

$$\sim N(\pi_t^*, \sigma_{i,t}^2 \equiv \sigma_{\pi,t}^2 + \sigma_{s,i,t}^2) \quad (12)$$

(where the second equation inserts assumption 4). The second signal, $x_{i,t}$, is distributed as

$$x_{i,t} \sim N(\pi_t^*, \sigma_{x,i,t}^2) \quad (13)$$

$s_{i,t}$ and $x_{i,t}$ are independent conditional on π_t^* .

$s_{i,t}$ represents the information that agents get from realized inflation, while $x_{i,t}$ represents the independent part of all the other information they observe, which might include policy announcements or information about other features of the economy. Corollary 1 then implies that the sensitivity of expectations to signals is

$$\frac{d}{ds_{i,t}} \kappa_{1,i,t} = \kappa_{2,i,t} \sigma_{i,t}^{-2} \quad (14)$$

The derivative represents how the agent's beliefs would change had they observed a different value of $s_{i,t}$, holding $x_{i,t}$ fixed. That result holds because the marginal sensitivity to *any* signal is equal to the posterior variance, $\kappa_{2,i,t}$, multiplied by the signal's precision.

Equation (14) can be applied to the data with the following result:

Proposition 1 *Under assumptions 2–5, the change in agent i 's expectations, $\Delta\kappa_{1,i,t}$, satisfies*

$$\Delta\kappa_{1,i,t} = \sigma_{i,t}^{-2}\kappa_{2,i,t}\hat{\pi}_{i,t} + \varepsilon_{i,t} \quad (15)$$

$$\text{where } \varepsilon_{i,t} \equiv \sigma_{i,t}^{-2}\kappa_{2,i,t}(s_{i,t} - \pi_t) + f_{i,t-1}(x_{i,t}) + o(s_{i,t} - \kappa_{1,i,t-1}) \quad (16)$$

$$\text{and } \hat{\pi}_{i,t} \equiv \pi_t - \kappa_{1,i,t-1} \quad (17)$$

$f_{i,t-1}(\cdot)$ is a function representing how expectations are updated given the unobservable component of agent i 's information, $x_{i,t}$.⁶

Changes in expectations naturally depend on the surprise in inflation, $\hat{\pi}_{i,t}$, and errors in agents' observation of inflation ($s_{i,t} - \pi_t$), scaled by their uncertainty and precision, based on equation (14). In addition, there is a residual term that accounts for other information that agents observe, $x_{i,t}$, via the function $f_{i,t-1}$, and higher order effects. The goal is to estimate equation (15) to evaluate whether the sensitivity of expectations to signals depends on uncertainty.

Proposition 15 shows that we can effectively think of a regression of the change in expectations on the product of uncertainty and the inflation surprise. Both the agents' observation errors and their other information will then be absorbed into the residual.

2.2 Bias in OLS estimates

In the empirical analysis, each agent's individual signal precision $\sigma_{i,t}^{-2}$ is not observable. Suppose, though, that $\kappa_{2,i,t}$ is observable. A natural feasible way to try to estimate the sensitivity of expectations to realized inflation – which the model says should be equal to $\sigma_{i,t}^{-2}$ – would be to regress $\Delta\kappa_{1,i,t}$ on $\kappa_{2,i,t}\hat{\pi}_{i,t}$. Unfortunately, that does not work in general.

Proposition 2 *The OLS estimator of the coefficient in a regression of $\Delta\kappa_{1,i,t}$ on $\kappa_{2,i,t}\hat{\pi}_{i,t}$, denoted $\hat{\beta}^{OLS}$, is*

$$\hat{\beta}^{OLS} = \widehat{\mathbb{E}}[\sigma_{i,t}^{-2}] + \widehat{\text{cov}}\left(\sigma_{i,t}^{-2}, \frac{\kappa_{2,i,t}\hat{\pi}_{i,t}^2}{\widehat{\mathbb{E}}[\kappa_{2,i,t}\hat{\pi}_{i,t}^2]}\right) + \frac{\widehat{\mathbb{E}}[\kappa_{2,i,t}\hat{\pi}_{i,t}\varepsilon_{i,t}]}{\widehat{\mathbb{E}}[\kappa_{2,i,t}\hat{\pi}_{i,t}^2]} \quad (18)$$

and $\widehat{\mathbb{E}}$ and $\widehat{\text{cov}}$ are the sample mean and covariance, respectively.

⁶Specifically, $f_{i,t-1}(x_{i,t}) = E[\pi_t^* | y_i^{t-1}, x_{i,t}, \hat{s}_{i,t} = 0] - E[\pi_{t-1}^* | y_i^{t-1}]$.

The first term in equation (18) is what we want to estimate. However, since information causes agents to update their beliefs, including their posterior variance, $\varepsilon_{i,t}$ will be correlated with $\kappa_{2,i,t}$, potentially leading to a bias because the third term can be nonzero. The second term in (18) shows that an additional bias can arise if signal precision is correlated with agents’ variances, $\kappa_{2,i,t}$. Those two variables naturally should be linked since agents’ variances depend on the information they’ve received.⁷ When precision is higher, all else equal, variance will fall. A valid estimator of $\widehat{\mathbb{E}}[\sigma_{i,t}^{-2}]$ needs to account for both sources of bias.

2.3 Instrumental variables methods


To avoid the endogeneity of uncertainty to both $\sigma_{i,t}^{-2}$ and $\varepsilon_{i,t}$, we instrument for uncertainty. Motivated by the results in [Kim and Binder \(2023\)](#), we use tenure in the survey as an instrument for uncertainty. The basic identifying assumption is then that tenure, $z_{i,t}$, is orthogonal to $\varepsilon_{i,t}$, while instrument relevance requires that $z_{i,t}$ is related to $\kappa_{2,i,t}$. [Kim and Binder \(2023\)](#) show (and we confirm below) that the latter condition holds. [Kim and Binder \(2023\)](#) discuss how an explanation for the decline in uncertainty associated with tenure is that when people are included in the survey they pay more attention to inflation news, and they might even learn from the survey itself. That suggests that people in the survey receive more precise signals than they did prior to the survey – $\sigma_{s,i,t}^2$ and $\sigma_{x,i,t}^2$ are lower – which then eventually leads them to have lower uncertainty. In other words, the “treatment” here is the survey itself.

Figure 3 visualizes the basic idea behind the identification. The x-axis represents the number of months since a respondent has entered the survey. The key assumption is that when people enter the survey, they pay more attention to inflation, which appears here as an increase in signal precision. That then causes their’ conditional variance ($\kappa_{2,i,t}$) to drift down. The identification in the regression comes from the within-person variation, where their uncertainty falls with tenure.

The figure also helps to see what would cause a bias, which is if attention varies systematically with tenure. If attention rises with tenure, that would counteract the effect of uncertainty and bias our estimates towards zero, whereas if it falls, that would compound the effect of uncertainty and cause us to overestimate the dependence of updating on uncertainty. Note that the assumption is not that signal precision is constant; it just needs to be unrelated to tenure (formalized in assumption 7 below). While attention obviously could systematically

⁷[Dew-Becker, Giglio and Molavi \(2026a\)](#) give general formulas for how uncertainty evolves as a function of signal precision, formalizing that intuition.

Figure 3: 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 ($\kappa_{2,i,t}$) 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.

vary with tenure, it is not obvious whether it would rise or fall, and on average across people it might be constant. Furthermore, even if there are differences across tenure, the first-order point is simply that being in the survey is very different from not being in it.

2.3.1 Problems with two-stage least squares

The usual approach to using an instrument is two-stage least squares (2SLS). Appendix B.2 discusses three reasons why 2SLS is ill-suited to this paper’s setting. The first major issue is that the endogenous variable here – $\kappa_{2,i,t}$ – enters the regression as an interaction. As Wooldridge (2015) discusses, that poses a problem for 2SLS because it is not valid to insert a fitted value from a first stage into an interaction. Instead, the interaction must be used as the dependent variable in the first stage, which can lead to a weak instruments problem that we show is practically relevant in our empirical application.

Beyond that, 2SLS estimates a weighted sum of $\sigma_{i,t}^{-2}$ across agents (the local average treatment effect). The weights in that average are not in general positive, though, meaning that the estimated coefficient from 2SLS here is not a conventional weighted average and it

is thus difficult or impossible to interpret economically.

2.3.2 Control function estimator

Wooldridge (2015) and Masten and Torgovitsky (2016) discuss how the control function approach to IV estimation is well suited to addressing both the problem of an average with negative weights and the endogenous variable entering as an interaction.

To use a control function type estimator, we need the following auxiliary assumption, analogous to a first-stage regression.

Assumption 6 *Given an instrument $z_{i,t}$ (tenure in our case), $\kappa_{2,i,t}$ has the first-stage representation*

$$\kappa_{2,i,t} = G(z_{i,t}, \hat{\pi}_{i,t}, v_{i,t}) \quad (19)$$

for a function G that is monotone in its third argument and where $v_{i,t}$ is a scalar random variable.

A special case of (19) is the usual linear first stage, $\kappa_{2,i,t} = b_0 + b_1 z_{i,t} + v_{i,t}$. Whatever is the functional form of G , the implication of assumption 6 is that $v_{i,t}$ can be recovered (up to a monotone transformation) and then used as a conditioning variable in the second stage.⁸

In addition to that, we need exogeneity and instrument relevance assumptions.

Assumption 7 *$\kappa_{2,i,t}$ is exogenous and uncorrelated with $\sigma_{i,t}^{-2}$ conditional on $v_{i,t}$ and $\hat{\pi}_{i,t}$. Specifically,*

$$\mathbb{E}[\sigma_{i,t}^{-2} \mid \kappa_{2,i,t}, v_{i,t}, \hat{\pi}_{i,t}] = \mathbb{E}[\sigma_{i,t}^{-2} \mid v_{i,t}, \hat{\pi}_{i,t}] \quad (20)$$

$$\mathbb{E}[\varepsilon_{i,t} \mid \kappa_{2,i,t}, v_{i,t}, \hat{\pi}_{i,t}] = \mathbb{E}[\varepsilon_{i,t} \mid v_{i,t}, \hat{\pi}_{i,t}] \quad (21)$$

where \mathbb{E} denotes the population expectation operator.

Assumption 7 requires that both the endogeneity of $\kappa_{2,i,t}$ and its correlation with the latent sensitivity, $\sigma_{i,t}^{-2}$, come entirely through $v_{i,t}$ and $\hat{\pi}_{i,t}$, as opposed to tenure, $z_{i,t}$. Equation (21) imposes the restriction that tenure is exogenous in the sense that the part of $\kappa_{2,i,t}$ driven by tenure (i.e. the part **not** driven by $v_{i,t}$ and $\hat{\pi}_{i,t}$) is unrelated to the residual $\varepsilon_{i,t}$. Equation (20) says that the part of $\kappa_{2,i,t}$ driven by tenure is also unrelated to signal precision in the cross-section, which is what we need in order to estimate the cross-sectional mean of $\sigma_{i,t}^{-2}$.

⁸Additional observable controls may be added to G if desired. What is critical is simply that, given the observables, G can be inverted to recover v .

That condition formalizes the restriction described above that we need tenure to be unrelated to signal precision.

Assumption 8 *The instrument relevance condition is*

$$\text{var}(\kappa_{2,i,t} \mid \hat{\pi}_{i,t}, v_{i,t}) \neq 0 \quad (22)$$

We then have the following standard result for the validity of the estimator.

Proposition 3 *Under assumptions 6–8, the second-stage regression can be estimated conditional on the values of $v_{i,t}$ and $\hat{\pi}_{i,t}$, with*

$$\Delta\kappa_{1,i,t} = \kappa_{2,i,t}\hat{\pi}_{i,t}\mathbb{E}[\sigma_{i,t}^{-2} \mid v_{i,t}, \hat{\pi}_{i,t}] + \mathbb{E}[\varepsilon_{i,t} \mid v_{i,t}, \hat{\pi}_{i,t}] + \eta_{i,t} \quad (23)$$

$$\text{where } \eta_{i,t} \equiv \Delta\kappa_{1,i,t} - \mathbb{E}[\Delta\kappa_{1,i,t} \mid \kappa_{2,i,t}, v_{i,t}, \hat{\pi}_{i,t}] \quad (24)$$

Assumption 8 ensures that there is variation in $\kappa_{2,i,t}\hat{\pi}_{i,t}$ so that $\mathbb{E}[\sigma_{i,t}^{-2} \mid v_{i,t}, \hat{\pi}_{i,t}]$ is identified.⁹

Equation (23) is simply a regression of $\Delta\kappa_{1,i,t}$ on $\kappa_{2,i,t}\hat{\pi}_{i,t}$ interacted with $\mathbb{E}[\sigma_{i,t}^{-2} \mid v_{i,t}, \hat{\pi}_{i,t}]$ along with whatever additional terms are necessary to capture $\mathbb{E}[\varepsilon_{i,t} \mid v_{i,t}, \hat{\pi}_{i,t}]$. These expectations are not directly observed, so Newey and Stouli (2021), for example, suggest fitting them with a sieve-type estimator in which the two conditional expectations are approximated by polynomials in $v_{i,t}$ and $\hat{\pi}_{i,t}$. For example, if the expectations are approximated as linear functions of $v_{i,t}$ and $\hat{\pi}_{i,t}$, then the regression would be of $\Delta\kappa_{1,i,t}$ on $[\kappa_{2,i,t}\hat{\pi}_{i,t}, 1] \otimes [1, v_{i,t}, \hat{\pi}_{i,t}]$ where \otimes denotes the Kronecker product. It is natural to also control for $v_{i,t}\hat{\pi}_{i,t}$ since that directly absorbs the potentially endogenous part of $\kappa_{2,i,t}\hat{\pi}_{i,t}$, so our baseline set of right-hand side variables will be $[\kappa_{2,i,t}\hat{\pi}_{i,t}, \hat{\pi}_{i,t}, 1] \otimes [1, v_{i,t}, \hat{\pi}_{i,t}]$.

The control function approach has the advantage that it yields estimates of $\mathbb{E}[\sigma_{i,t}^{-2} \mid v_{i,t}, \hat{\pi}_{i,t}]$. To get the unconditional mean of $\sigma_{i,t}^{-2}$, which is the parameter of interest, we can simply average across values of $v_{i,t}$ and $\hat{\pi}_{i,t}$. As Masten and Torgovitsky (2016) emphasize, the control function (CF) approach therefore yields an estimate of the parameter of interest, $\mathbb{E}[\sigma_{i,t}^{-2}]$, rather than a weighted average.

The CF approach therefore addresses all three concerns in 2SLS estimation: the first-stage does not involve an interaction, thus reducing the weak instruments problem; even if $\sigma_{i,t}^{-2}$ is correlated with $\kappa_{2,i,t}$ the method can still yield an estimate of its unconditional mean; and since the estimator is not a weighted average, there is no concern of the weights potentially being negative.

⁹To derive equation (23), take the expectation of equation (15) conditional on $\{\kappa_{2,i,t}, v_{i,t}, \hat{\pi}_{i,t}\}$ and apply assumption 7.

2.4 Estimation for squared changes in beliefs

A simpler approach to estimating the sensitivity of expectations to news is just to look at the magnitude of their changes. If we take the benchmark case from section 1 (i.e. assumption 2) that agents have a single composite signal $y_{i,t}$,¹⁰ then equation (4) yields the following:

Proposition 4 *Under assumptions 2 and 3,*

$$\Delta\kappa_{1,i,t}^2 = \kappa_{2,i,t}^2 \sigma_{i,t}^{-4} \hat{y}_{i,t}^2 + \lambda_{i,t} \quad (25)$$

where $\lambda_{i,t}$ is a residual.

Note that proposition 4 does not require assumptions 4 and 5; all it needs is the basic assumptions that yield the power series for the conditional expectation from theorem 1.

Now consider a control-function type regression – using the same control function, (19), as above – for $(\Delta\kappa_{1,i,t})^2$. The analog to equation (23) in proposition 3 is

$$\Delta\kappa_{1,i,t}^2 = \kappa_{2,i,t}^2 \mathbb{E} [\sigma_{i,t}^{-4} \hat{y}_{i,t}^2 | v_{i,t}] + \mathbb{E} [\lambda_{i,t} | v_{i,t}] + \mu_{i,t} \quad (26)$$

where $\mu_{i,t}$ is a residual.¹¹ That implies a regression of $\Delta\kappa_{1,i,t}^2$ on $\kappa_{2,i,t}^2$ interacted with a constant and functions of the control function $v_{i,t}$.

One assumption that allows for a simple interpretation of the coefficient in (26) is that $\hat{y}_{i,t}^2$ is unpredictable by $v_{i,t}$.¹² Then

$$\mathbb{E} [\sigma_{i,t}^{-4} \hat{y}_{i,t}^2 | v_{i,t}] \approx \mathbb{E} [\sigma_{i,t}^{-2} | v_{i,t}] \quad (30)$$

¹⁰In the context of the analysis of inflation as a signal (which uses stronger assumptions than the present results), $y_{i,t}$ would be constructed as $y_{i,t} = (\sigma_{x,i,t}^{-2} x_{i,t} + \sigma_{s,i,t}^{-2} s_{i,t}) (\sigma_{x,i,t}^{-2} + \sigma_{s,i,t}^{-2})^{-1}$, but this section does not need the particular structure from assumptions 4 and 5.

¹¹Formally, assumptions 6 and 8 are still required and the analog to assumption 7 is

$$\mathbb{E} [\sigma_{i,t}^{-4} | \kappa_{2,i,t}, v_{i,t}, \hat{\pi}_{i,t}] = \mathbb{E} [\sigma_{i,t}^{-4} | v_{i,t}, \hat{\pi}_{i,t}] \quad (27)$$

$$\mathbb{E} [\lambda_{i,t} | \kappa_{2,i,t}, v_{i,t}, \hat{\pi}_{i,t}] = \mathbb{E} [\lambda_{i,t} | v_{i,t}, \hat{\pi}_{i,t}] \quad (28)$$

¹²Using theorem 2 and defining $\kappa_{n,i,t}^0$ to be the n th posterior cumulant conditional on $\hat{y}_{i,t} = 0$, we have

$$\kappa_{2,i,t} - \kappa_{2,i,t}^0 = \kappa_{3,i,t} \sigma_{i,t}^{-2} \hat{y}_{i,t} + \frac{1}{2} \kappa_{4,i,t} (\sigma_{i,t}^{-2} \hat{y}_{i,t})^2 + o(\hat{y}_{i,t}^2) \quad (29)$$

The relationship between $\kappa_{2,i,t}$, and hence $v_{i,t}$, and $\hat{y}_{i,t}^2$ therefore comes through the second-order effect of signals on uncertainty, which in a rough sense implies the relationship might be weak. Additionally, if agents' beliefs are approximately Gaussian, then $\kappa_{4,i,t} \approx 0$, giving another reason the relationship could be weak.

under the approximation that $E[\hat{y}_{i,t}^2 | \sigma_{i,t}^{-2}, v_{i,t}] \approx E[\sigma_{i,t}^2 | v_{i,t}]$. The coefficient on $\kappa_{2,i,t}^2$ will then again be an estimate of $E[\sigma_{i,t}^{-2}]$. More generally, though, if that assumption is violated, the regression still gives an estimate of the structural relationship between changes in uncertainty and the cross-sectional average magnitude of changes in beliefs.

3 Panel analysis

3.1 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.


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: “Now we would like you to think about the different things that may happen to inflation over the next 12 months... what would you say is the percent chance that, over the next 12 months...” followed by a list of bins for inflation: >12%, 8-12%, etc. Equivalent questions are asked for annual inflation starting 24 and 48 months in the future. We use both the point estimates and the bins.

The bins are critical because they make it possible to calculate a mean and conditional variance for each person. 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 each respondent’s reported bin probabilities.¹³


The paper’s sample period is from the beginning of the SCE in June, 2013, through July, 2025. To get some context, figure 4 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.

¹³They 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 4: 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.

3.2 Calculating inflation surprises

We calculate the inflation surprise in each month for each person as realized inflation – PCE or CPI inflation, headline or core – minus that person’s expected one-year inflation – either the point forecast or the mean from the bins – in the previous month.

Since the change in expectations, $\Delta\kappa_{1,i,t}$ is on the left-hand side of the regression, there might be a concern about using $\kappa_{1,i,t-1}$ on the right-hand side of the regression in constructing the inflation surprise (e.g. if there is measurement error in the expectations). For that reason, 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. That is, for those cases $\hat{\pi}_{i,t} = \pi_t - \kappa_{1,i,t-2}$.

In what follows, $\hat{\pi}_{i,t}$ continues to represent the generic inflation surprise, which may be measured in any of the ways described above – there are 16 total possible permutations.¹⁴

In the main text, inflation is measured as the CPI headline value. The dependent variable, $\Delta\kappa_{1,i,t}$ is measured based on the mean from respondents’ probability distributions. Inflation surprises are then measured as realized CPI headline inflation minus lagged expectations. The main text reports results calculating the surprise relative to either the bin mean or the point forecast. Robustness to all these choices is reported in the appendix.

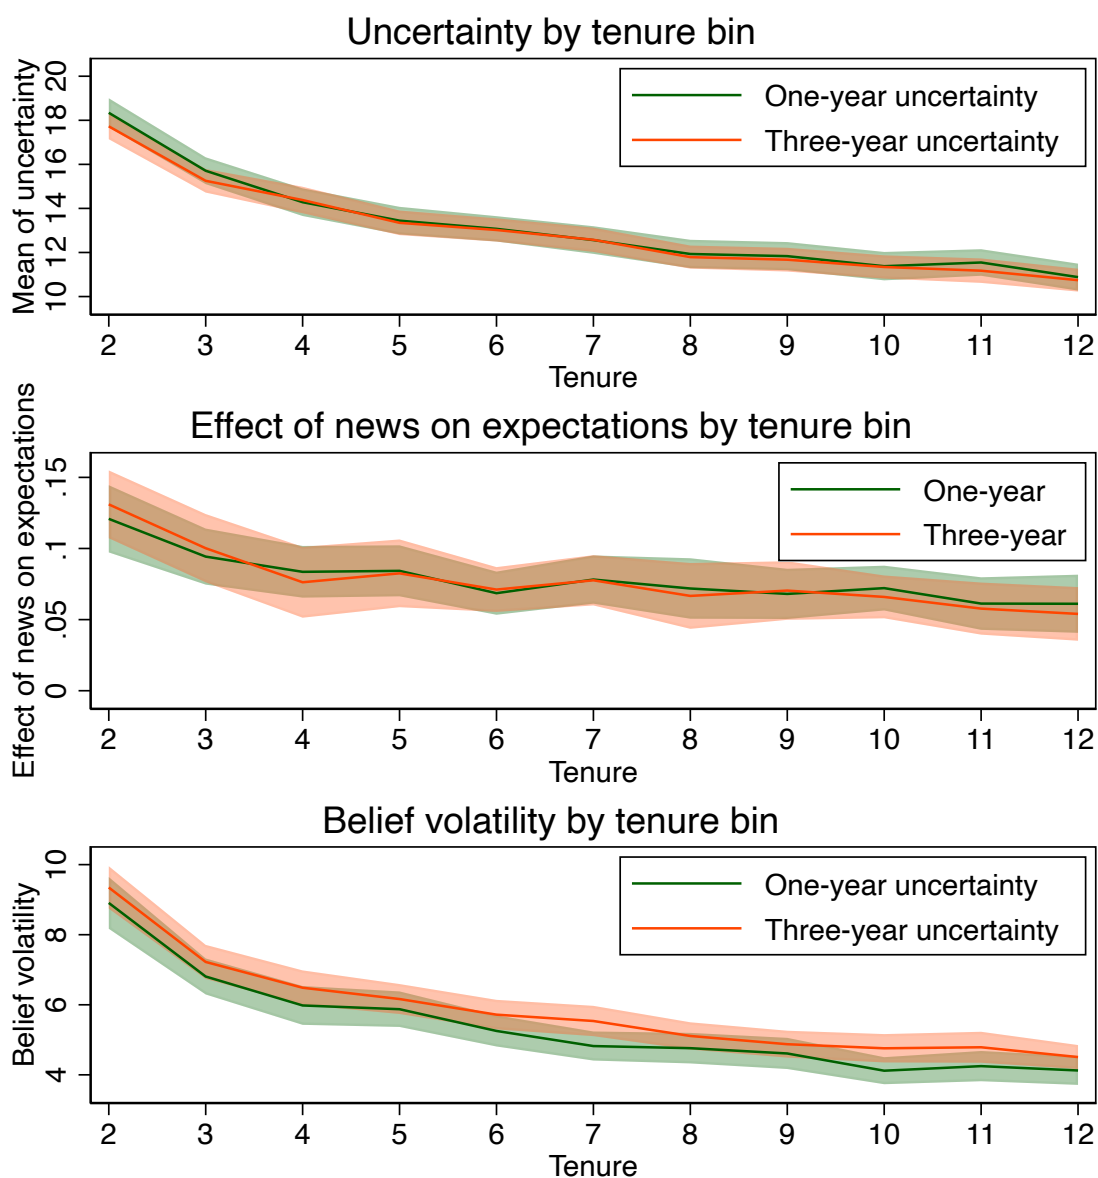
3.3 First stage

The first stage, following [Kim and Binder \(2023\)](#), uses tenure in the survey as an instrument for uncertainty. The top panel of figure 5 plots average uncertainty (the variance, $\kappa_{2,i,t}$) at both the one- and three-year horizons by tenure in the survey. There is a clear negative relationship, which is statistically well estimated. Most of the decline is in the first few months of tenure, but even between months 7 and 12 there is a decline in average uncertainty of 14 percent.

For much of the analysis, we use a simple linear first stage (i.e. the function G in (19) is linear). Table 1 reports results for the first stage in the CF approach. 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.

¹⁴The 16 permutations are from the cross of four binary choices: (1) PCE or CPI; (2) headline or core; (3) one- or three-year horizon; (4) point estimate or mean from the bins.

Figure 5: 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.

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 bin mean or point forecast – with bins, again, the independent variable is entered with two lags since it appears on the right-hand side in the baseline second stage. In columns 3–4 the t-statistic on tenure shrinks somewhat, but not enough to raise any weak identification concerns. Looking at the coefficients themselves, 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.

Table 1: First stage regressions

| (a) First Stage, 1 year inflation | | | | |
|-----------------------------------|-----------------------|-----------------------|----------------------|----------------------|
| Dep var: | (1) Unc. | (2) Unc. | (3) Unc. | (4) Unc. |
| Tenure | -0.377*** (-16.71) | -0.264*** (-11.17) | -0.183*** (-9.56) | -0.195*** (-8.24) |
| Infl. Surpr. | | | -0.602*** (-7.63) | 0.177*** (3.50) |
| 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 | | | | |
|-----------------------------------|-----------------------|-----------------------|----------------------|----------------------|
| Dep var: | (1) Unc. | (2) Unc. | (3) Unc. | (4) Unc. |
| Tenure | -0.393*** (-17.16) | -0.294*** (-12.31) | -0.183*** (-8.89) | -0.231*** (-9.84) |
| Infl. Surpr. | | | -0.601*** (-7.65) | 0.271*** (5.61) |
| 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. Robust t-statistics allowing for two-way clustering by person and time are reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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) cutoff for this case is no smaller than 12, indicating that 2SLS is close to weakly identified here.

3.4 Reduced-form estimates

3.4.1 News regressions

Table 2, columns (1) and (2), reports results for reduced-form regressions of $\Delta\kappa_{1,i,t}$ on $\hat{\pi}_{i,t}$ and $\hat{\pi}_{i,t}z_{i,t}$ for the two inflation surprise measures (point estimate and bin mean). Panel (a) uses 1-year expectations, panel (b) 3-year expectations. Across both columns, the coefficient on the inflation surprise is positive and the coefficient on $\hat{\pi}_{i,t}z_{i,t}$ is negative, both as expected. Inflation expectations rise when inflation is surprisingly high, and the slope of that relationship decreases with tenure.

Table 2: Reduced form regressions

| (a) 1 year inflation | | | | |
|------------------------------|-----------------------------|-----------------------|---------------------------------|----------------------|
| | Dep. var.: $\Delta\kappa_1$ | | Dep. var.: $(\Delta\kappa_1)^2$ | |
| | (1) | (2) | | (3) |
| Tenure \times Infl. Surpr. | -0.0046*** (-8.13) | -0.0021*** (-3.78) | Tenure | -0.40*** (-19.40) |
| Infl. Surpr. | 0.1112*** (12.03) | 0.0445*** (7.31) | | |
| Type of infl. surpr. | Point | Mean | | |
| R^2 | 0.0270 | 0.0040 | R^2 | 0.01 |
| Obs. | 103,888 | 93,460 | Obs. | 118,141 |

| (b) 3 year inflation | | | | |
|------------------------------|-----------------------------|---------------------|---------------------------------|----------------------|
| | Dep. var.: $\Delta\kappa_1$ | | Dep. var.: $(\Delta\kappa_1)^2$ | |
| | (1) | (2) | | (3) |
| Tenure \times Infl. Surpr. | -0.0025*** (-4.31) | -0.0007 (-1.28) | Tenure | -0.40*** (-21.73) |
| Infl. Surpr. | 0.0611*** (8.59) | 0.0244*** (4.06) | | |
| Type of infl. surpr. | Point | Mean | | |
| R^2 | 0.0070 | 0.0010 | R^2 | 0.01 |
| Obs. | 101,438 | 92,765 | Obs. | 118,620 |

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

To get a sense of magnitudes, the average sensitivity of one-year expectations to news at $tenure = 2$ in columns (1) and (2) is 0.10 and 0.06, compared to 0.04 and 0.02 at 12 months. That is, a person at the end of their time in the sample is 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 5 plots estimates from regressions of $\Delta\kappa_{1,i,t}$ 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 fall by about half for agents at the end of their tenure compared to the beginning.

Kim and Binder (2023) show that in addition to uncertainty declining with tenure, average inflation expectations do, also. That should cause the average of $\Delta\kappa_{1,i,t}$ 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 2 but including tenure as an additional control. The coefficient on the interaction $\hat{\pi}_{i,t}z_{i,t}$ is essentially unaffected, which is not surprising. 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.

3.4.2 Volatility in expectations

The third column of table 2 reports results from regressions of $(\Delta\kappa_{1,i,t})^2$ 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 5 plots the average of $(\Delta\kappa_{1,i,t})^2$ by tenure bins (equivalent to a regression with tenure dummies). 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.

3.5 Second-stage estimates

3.5.1 News regressions

Table 3 reports results from the second stage of the control function estimator.¹⁵ The first row reports the mean sensitivity of expectations to $\kappa_{2,i,t}\hat{\pi}_{i,t}$ (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

¹⁵Standard 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 3 account for uncertainty in the first stage. As above, standard errors are two-way clustered by person and time.

its persistent component and also the difference between what agents “observe” and actual realized inflation.

Table 3: Second Stage Regression

| | 1 year inflation | | 3 year inflation | |
|--|-----------------------|-----------------------|-----------------------|---------------------|
| | (1) | (2) | (3) | (4) |
| Unc. \times Infl. Surpr., avg marginal effect | 0.0229*** (6.83) | 0.0114*** (3.44) | 0.0120*** (3.71) | 0.0032 (1.24) |
| $v_{i,t} \times$ Infl. Surpr., avg marginal effect | -0.0245*** (-7.30) | -0.0127*** (-3.88) | -0.0134*** (-4.12) | -0.0046* (-1.73) |
| 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 $\Delta\kappa_{1,i,t}$ (measured based on the bin mean) on interactions of uncertainty (Unc. = $\kappa_{2,i,t}$), inflation surprise ($\hat{\pi}_{i,t}$), and control function ($v_{i,t}$). The regression includes the terms obtained from the Kroneker product of (inflation surprise; uncertainty \times inflation surprise) with (a constant, the control function component, and inflation surprise). All the coefficients of this regression are reported in table A.4. This table reports the mean sensitivity of expectations to $\kappa_{2,i,t}\hat{\pi}_{i,t}$ in the first row and the mean sensitivity of expectations to $v_{i,t}\hat{\pi}_{i,t}$ (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. t -statistics in parentheses 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 3 reports the mean sensitivity of expectations to $v_{i,t}\hat{\pi}_{i,t}$, 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 $v_{i,t}$ 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 $\kappa_{2,i,t}\hat{\pi}_{i,t}$, they are broadly consistent with each other and uniformly positive and significant. Table A.6 replicates table 3 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 – mostly importantly, higher powers of various variables (recalling the sieve interpretation) – and finds similar results.

For the sake of completeness, table A.8 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 may not be particularly important.

3.5.2 Volatility in expectations

Table 4 reports results for the second-stage regression of $(\Delta\kappa_{1,i,t})^2$ on uncertainty. The estimated sensitivity to $\kappa_{2,i,t}^2$ is, under the conditions discussed above, again an estimate of $E[\sigma_{i,t}^{-2}]$, 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 4: Second Stage Regression for $(\Delta\kappa_1)^2$

| | 1-year inflation | | 3-year inflation | |
|-----------------------------------|---------------------|----------------------|---------------------|----------------------|
| | (1) | (2) | (3) | (4) |
| $(Unc.)^2$, avg marginal effect | 0.1768*** (7.25) | 0.0942*** (13.49) | 0.2483*** (8.18) | 0.0862*** (13.80) |
| $v_{i,t}^2$, avg marginal effect | 0.0216** (2.27) | 0.0121*** (4.68) | 0.0543*** (4.83) | 0.0207*** (9.03) |
| Type of inflation expectation | Point | Mean | Point | Mean |
| Observations | 103,371 | 118,141 | 102,113 | 118,620 |

Note: Results from second-stage regressions of $(\Delta\kappa_{1,i,t})^2$ on interactions of uncertainty and the control-function component. Uncertainty is $Unc_{i,t} = \kappa_{2,i,t}$. The regression includes terms obtained from the Kronecker product $(1, Unc_{i,t}^2) \otimes (1, v_{i,t}, v_{i,t}^2, v_{i,t}^{inv})$, where $v_{i,t}$ is demeaned and $v_{i,t}^{inv} = (v_{i,t} - \min(v_{i,t}) + 10)^{-1}$. All coefficients from the second-stage regression are reported in the appendix. This table reports the average marginal effect of $Unc_{i,t}^2$ and the average marginal effect 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). t -statistics in parentheses 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 $E[\sigma_{i,t}^{-2}]$ are significantly larger in table 4 than in table 3. That is exactly what we should expect – agents receive many signals about future inflation beyond realized current inflation. The results in table 3 indicated that the signal agents got from inflation had an error with effective standard deviation of 7–18%. Table 4 implies that the standard deviation of the effective error in their *overall* signal is between 2 and 3.4%.

The second row of table 4 reports the sensitivity to $v_{i,t}^2$. In this case, in contrast to table 3, 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 3, the negative coefficient on $v_{i,t}\hat{\pi}_{i,t}$ 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 4 is positive.

4 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.

The theoretical analysis stressed the endogeneity of uncertainty to signal precision. Taking a cross-sectional average helps alleviate that concern because while it seems very likely that there are people with high uncertainty because they pay little attention to inflation, it seems less likely that variation over time in average uncertainty is driven by variation over time in the precision of signals agents observe. It is certainly not *impossible*, just not as prominent an issue. Furthermore, in this case we are not so much trying to test the model as simply ask whether uncertainty is useful for predicting sensitivity to shocks, which, again, is a forecasting problem and not a causal identification problem.

To align with typical aggregate time series, we first collapse the panel structure into cross-sectional medians in each period. We then define inflation news as inflation in month t minus the cross-sectional median of 12-month expected inflation 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 4 reports results from four regressions. The first two are of the change in 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 $t - 1$ ($\kappa_{2,i,t-1}$). We subtract the time-series mean from uncertainty so that the interaction coefficient has a forecasting interpretation: it gives a forecast at date $t - 1$ of the sensitivity of expectations to inflation surprises on date t .

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

Table 5: Aggregate Analysis

| | (1) | (2) | (3) | (4) |
|-------------------------------|------------------------|------------------------|----------------------------|----------------------------|
| Dep. var.: | $\Delta\kappa_1$ (1yr) | $\Delta\kappa_1$ (3yr) | $(\Delta\kappa_1)^2$ (1yr) | $(\Delta\kappa_1)^2$ (3yr) |
| Infl. Surprise | 0.01931** (2.40) | 0.01146*** (3.49) | | |
| Lagged unc. | -0.01638 (-1.02) | -0.01846 (-1.28) | | |
| Surprise \times Lagged unc. | 0.00870*** (3.42) | 0.00456** (2.09) | | |
| Lagged (unc.) ² | | | 0.00372*** (5.09) | 0.00197*** (3.72) |
| Constant | 0.00757 (0.38) | 0.00385 (0.39) | 0.03494*** (6.44) | 0.01735*** (6.94) |
| 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. All variables on the right hand side are time-series demeaned. Newey-West t -statistics with 12 lags are reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.0, respectively. In the first column, the coefficient on news is 0.02 – a one percentage point surprise in annualized monthly inflation raises one-year expectations by 0.02 percentage points. The coefficient on the interaction with uncertainty is 0.009, meaning that the standard deviation of that sensitivity is 0.013. For three-year expectations, the mean sensitivity to inflation surprises is, naturally, smaller at 0.01, and the interaction with uncertainty is 0.005.

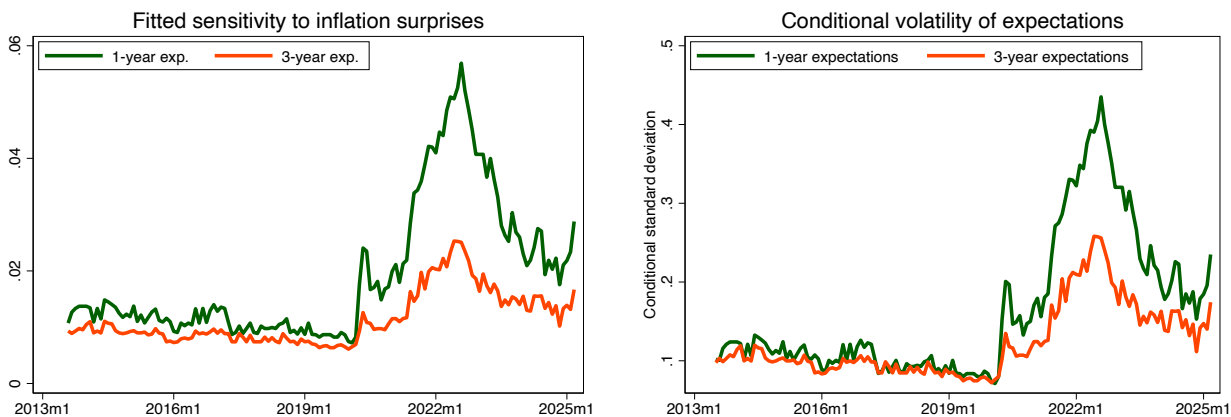
The left-hand panel of figure 6 plots the fitted sensitivity of one- and three-year expectations from the two regressions. There is a striking degree of variation. At its peak in 2022, the sensitivity of one-year expectations is estimated to be 0.06. Exactly when inflation was highest, expectations were 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 late 2024, sensitivity remains nearly twice as high as it was then.

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 two columns report the regression of squared changes in median expectations on (demeaned) lagged squared uncertainty. The coefficients are again highly statistically significant. To see the effects on the predicted volatility, the right-hand panel of figure 6 plots the square roots of the fitted values from these two regressions. These values represent conditional standard deviations for one- and three-year median inflation expectations. 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 at the end of 2024 remained more than twice as high as pre-covid.

Figure 6: 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.

5 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 three 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.

- Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston**, “Anchored inflation expectations,” *American Economic Journal: Macroeconomics*, 2023, 15 (1), 1–47.
- Coibion, Olivier and Yuriy Gorodnichenko**, “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 are Meaningful and Important,” *Annual Review of Economics*, 2024.
- Dew-Becker, Ian, Stefano Giglio, and Pooya Molavi**, “Learning and the emergence of nonlinearity in financial markets,” 2026. Working paper.
- , –, and –, “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.
- 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.
- 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.
- 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.
- 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.
- 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,” *Journal of Human Resources*, 2015, *50* (2), 420–445.

A Proofs

A.1 Derivation of theorems 1 and 2

Consider an arbitrary random variable x with prior distribution p . There is a signal $y \sim N(x, \sigma^2)$. The posterior cumulant generating function of x given $y = Y$ is

$$CGF_{x|Y}(s) = \log E[\exp(sx) | y = Y] \quad (\text{A.1})$$

$$= \log \frac{\int \exp(sx) p(x) \exp\left(-\frac{1}{2}(Y-x)^2 \sigma^{-2}\right) dx}{\int p(x) \exp\left(-\frac{1}{2}(Y-x)^2 \sigma^{-2}\right) dx} \quad (\text{A.2})$$

$$= \log \int \exp(sx) p(x) \exp\left(-\frac{1}{2}(Y-x)^2 \sigma^{-2}\right) dx - \log \int p(x) \exp\left(-\frac{1}{2}(Y-x)^2 \sigma^{-2}\right) dx \quad (\text{A.3})$$

(note here that the CGF always exists in a neighborhood of zero, even if the prior moments of x fail to exist, because the posterior density is scaled by $\exp(-x^2\sigma^{-2}/2)$, meaning that its tails decay sufficiently fast that all of its moments exist).

Now differentiate $CGF_{x|Y}(s)$ with respect to s

$$\frac{d}{ds} CGF_{x|Y}(s) = \frac{d}{ds} \left[\log \int \exp(sx) p(x) \exp\left(-\frac{1}{2}(Y-x)^2 \sigma^{-2}\right) dx \right] \quad (\text{A.4})$$

$$= \frac{\int x \exp((s + Y\sigma^{-2})x) p(x) \exp\left(-\frac{1}{2}x^2\sigma^{-2}\right) dx}{\int \exp((s + Y\sigma^{-2})x) p(x) \exp\left(-\frac{1}{2}x^2\sigma^{-2}\right) dx} \quad (\text{A.5})$$

Evaluated at $s = 0$, $\frac{d}{ds} CGF_{x|Y}(s)$ is the first posterior cumulant of x - i.e. the posterior mean. In addition, note that all derivatives of $\frac{d}{ds} CGF_{x|Y}(s)$ with respect to Y/σ^2 are the same as those with respect to s . More generally, for all $j + k > 1$

$$\frac{d^j}{d(Y\sigma^{-2})^j} \frac{d^k}{ds^k} CGF_{x|Y}(s) = \frac{d^{j+k}}{ds^{j+k}} CGF_{x|Y}(s) \quad (\text{A.6})$$

Since $\frac{d^k}{ds^k} CGF_{x|Y}(s) = \kappa_k(x | y = Y)$, we have

$$\frac{d^j}{d(Y\sigma^{-2})^j} \kappa_k(x | y = Y) = \kappa_{j+k}(x | y = Y) \quad (\text{A.7})$$

which gives a power series representation for $E[x | y = Y]$ around $Y = a$ for an arbitrary a by setting k above to 1, yielding theorem 1. Theorem 2 follows from (A.7) by setting $k = 2$

and $j = 1$.

A.2 Proof of proposition 1

First, define the expectation conditional on the realized signal being equal to the prior mean,

$$\kappa_{1,i,t}^0 \equiv \mathbb{E} [\pi_t^* | x_i^t, s_i^{t-1}, \hat{s}_{i,t} = 0] \quad (\text{A.8})$$

$$\text{where } \hat{s}_{i,t} = s_{i,t} - \kappa_{1,i,t-1} \quad (\text{A.9})$$

$\kappa_{1,i,t}^0$ represents the value that expectations would have taken at the end of period t – conditional on x_i^t and s_i^t – had the surprise in the signal, $\hat{s}_{i,t}$, been equal to zero. We can get to $\kappa_{1,i,t}^0$ either as an update from $t - 1$ or by using theorem 1 to go backwards from $\kappa_{1,i,t}$. For the latter, from (14), we have

$$\kappa_{1,i,t}^0 = \kappa_{1,i,t} - \kappa_{2,i,t} \sigma_{i,t}^{-2} \hat{s}_{i,t} - o(\hat{s}_{i,t}) \quad (\text{A.10})$$

where the $o(\hat{s}_{i,t})$ term involves higher-order powers of $\hat{s}_{i,t}$. For the forward update from $t - 1$,

$$\kappa_{1,i,t}^0 = \kappa_{1,i,t-1} + f_{i,t-1}(x_{i,t}) \quad (\text{A.11})$$

where $f_{i,t-1}$ is a function representing how expectations are updated given the unobservable component of agents' information.¹

The final issue is that we do not observe the individual-specific signals, $s_{i,t}$. Instead, we can at best observe $\hat{\pi}_{i,t} \equiv \pi_t - \kappa_{1,i,t-1}$. Combining that with equations (A.10) and (A.11) yields the result.

B Extensions and additional discussion

B.1 Some deviations from fully rational Bayesianism

The results here do not require that people are perfect Bayesians, at least in the usual sense of combining Bayes' theorem with rational expectations. [Dew-Becker, Giglio and Molavi \(2026a\)](#) also discuss this point, and [Molavi \(2025\)](#) analyzes it in depth. In particular, assumptions 2 and 3 do not impose any of the following:

1. That agents use the correct 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)

¹Specifically, $f_{i,t-1}(x_{i,t}) = E[\pi_t^* | y_i^{t-1}, x_{i,t}, \hat{s}_{i,t} = 0] - E[\pi_{t-1}^* | y_i^{t-1}]$.

2. That agents’ prior is in any sense “correct”. 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.
3. That the signals agents observe are actually Gaussian conditional on future inflation.
4. That the agents incorporate all information they receive. They could irrationally or inefficiently ignore some information, and unreasonably privilege other sources.
5. That the agents properly weight all information they receive. For example, agents might receive many Gaussian signals, which can be combined into a single value and used for updating. It is possible that they do that combination incorrectly.
6. That agents all receive the same information or share priors. For example, their beliefs could be affected by different life histories, as in [Malmendier and Nagel \(2016\)](#).

What is important here is not actually that agents are true Bayesians. Theorem 1 simply requires that they use an updating rule that has the same algebraic structure as Bayes’ rule. Obviously people might not actually do that, in which case the model’s predictions should fail in the data. The claim is not that every possible belief updating rule fits into the structure of theorem 1, and in fact the goal of the next sections of the paper is to evaluate whether theorem 1 has any descriptive power.

B.2 Issues with two-stage least squares

The usual approach to instrumental variables estimation is two-stage least squares. This section explains why that method is not well suited to this paper’s setting.

Since the independent variable in the regression is $\kappa_{2,i,t}\hat{\pi}_{i,t}$, the natural instrument is $z_{i,t}\hat{\pi}_{i,t}$. The 2SLS estimated coefficient is then

$$\hat{\beta}^{2SLS} = \hat{\mathbb{E}} \left[\frac{\widetilde{z_{i,t}\hat{\pi}_{i,t}\kappa_{2,i,t}\hat{\pi}_{i,t}}}{\hat{E} \left[\widetilde{z_{i,t}\hat{\pi}_{i,t}\kappa_{2,i,t}\hat{\pi}_{i,t}} \right]} \sigma_{i,t}^{-2} \right] + \frac{\hat{\mathbb{E}} \left[\widetilde{z_{i,t}\hat{\pi}_{i,t}\tilde{\varepsilon}_{i,t}} \right]}{\hat{\mathbb{E}} \left[\widetilde{z_{i,t}\hat{\pi}_{i,t}\kappa_{2,i,t}\hat{\pi}_{i,t}} \right]}$$

where for any variable x , \tilde{x} represents the demeaned value. That equation leads to two concerns. The first is that $\hat{\beta}^{2SLS}$, even in the case where $\mathbb{E} \left[\widetilde{z_{i,t}\hat{\pi}_{i,t}\varepsilon_{i,t}} \right] = 0$, does not in general estimate the mean of $\sigma_{i,t}^{-2}$. Instead, as usual for 2SLS, it estimates a type of local

average treatment effect. Since we have a continuous instrument, we get the standard result that the estimated coefficient is a weighted average of the $\sigma_{i,t}^{-2}$. That average will be tilted towards the observations for which $\widetilde{z_{i,t}\hat{\pi}_{i,t}\kappa_{2,i,t}\hat{\pi}_{i,t}}$ is relatively high, which is similar to the bias in OLS above.

More troublingly, though, without strong additional assumptions, there is no reason to necessarily think that the weights $\frac{\widetilde{z_{i,t}\hat{\pi}_{i,t}\kappa_{2,i,t}\hat{\pi}_{i,t}}}{\mathbb{E}[\widetilde{z_{i,t}\hat{\pi}_{i,t}\kappa_{2,i,t}\hat{\pi}_{i,t}}]}$ are all positive. They will only be uniformly positive if it is the case that, observation-by-observation, whenever $\widetilde{\kappa_{2,i,t}\hat{\pi}_{i,t}} \geq 0$, we also have $\widetilde{z_{i,t}\hat{\pi}_{i,t}} \geq 0$ or $\widetilde{z_{i,t}\hat{\pi}_{i,t}} \leq 0$ (i.e. they always have identical signs or always have opposite signs), which is obviously a strong restriction. Without that, nothing ensures that the estimated coefficient is actually an average of the values of $\sigma_{i,t}^{-2}$ in the usual sense.

Finally, there is also good reason to be concerned that the first stage might be weak here, because the instrument, $z_{i,t}$, is multiplied by $\hat{\pi}_{i,t}$. In particular, the economically natural first stage is a regression of $\kappa_{2,i,t}$ on $z_{i,t}$. Section B.2 gives an informal argument that the F -statistic in the weighted regression should, on average, be smaller than that in the unweighted regression by a factor equal to the kurtosis of $\hat{\pi}_{i,t}$. If $\hat{\pi}_{i,t}$ is normally distributed, for example, then its kurtosis is equal to 3 (which is close to what we observe in the data). The first-stage F -statistic in the weighted regression in that case will be smaller by a factor of 3 than in the unweighted regression, which can be expected to potentially lead to a weak-instrument problem.

As Wooldridge (2015) discusses, it is not in general valid to estimate a first stage regression of κ_2 on z and then insert the fitted value as an interaction in the second stage. But alternative instrumental variables methods do allow us to use the uninteracted first stage.

B.3 Estimation example

This section presents a simple example with distributions for the various terms in the analysis for which OLS estimates of the relationship between information sensitivity and uncertainty are biased towards zero, while a control function approach can (in population) perfectly recover the structural parameters. The example also helps motivate the inclusion of the inverse of the estimated control function in the empirical specification. In this section we ignore the time component of the model and simply index observations by i .

Suppose there is a random variable

$$s_i \sim U(\delta, 1) \tag{A.12}$$

s_i will represent an idiosyncratic component of signal precision. In addition, $\hat{\pi}_i$ and ε_i have arbitrary mean-zero distributions, and s_i , $\hat{\pi}_i$, ε_i are jointly independent. Finally, there is an instrument z_i , which can be thought of as tenure, though we leave it unstructured here. The scale of the various variables will be essentially irrelevant. The only functional form that will matter is that for s_i , which will allow for some closed-form calculations.

For uncertainty and signal precision we assume

$$\sigma_i^{-2} = a + bs_i \tag{A.13}$$

$$\kappa_{2,i} = s_i^{-1/2} + z_i \tag{A.14}$$

Finally, we assume that the theoretical structure from the paper holds in the sense that

$$\Delta\kappa_{1,i} = \sigma_i^{-2}\kappa_{2,i}\hat{\pi}_i + \varepsilon_i \tag{A.15}$$

B.3.1 Motivation

It is worth briefly motivating the functional forms above, specifically for σ_i^{-2} and $\kappa_{2,i}$. Again, s_i is meant to capture cross-sectional variation in signal precision or, equivalently, the attention agents pay to inflation. Using results from [Dew-Becker et al. \(2026b\)](#), the steady-state level of uncertainty given signal precision of s is, roughly,

$$\kappa_2^{SS} = \frac{\mathbb{E} [d(\pi_t^*)^2]^{1/2}}{s^{1/2}} \tag{A.16}$$

where $E[d(\pi_t^*)^2]$ is the expected change in squared change in trend inflation. The equation comes from solving the usual Riccati equation that appears in filtering. That motivates the inverse relationship between $\kappa_{2,i}$ and s_i . z_i represents the variation in uncertainty due to tenure.

Certainly this example is stylized. It does not formalize the idea that z_i should be related to σ_i^{-2} across agents, but not within agents over time.

B.3.2 Control function regressions

The assumption in the main text is that $\kappa_{2,i} = G(z_i, \hat{\pi}_i, v_i)$. Here, $\kappa_{2,i}$ has such a form. G is linear, $\hat{\pi}_i$ does not matter, and, most importantly, the control function we will recover is $v_i = s_i^{-1/2}$.

The identification conditions here from assumption 7 are

$$\begin{aligned}\mathbb{E}[\sigma_i^{-2} | \kappa_{2,i}, v_i, \hat{\pi}_i] &= \mathbb{E}[\sigma_i^{-2} | v_i, \hat{\pi}_i] \\ \mathbb{E}[\varepsilon_i | \kappa_{2,i}, v_i, \hat{\pi}_i] &= \mathbb{E}[\varepsilon_i | v_i, \hat{\pi}_i]\end{aligned}$$

The first condition holds since σ_i^{-2} is a function of v_i alone. The second holds trivially since ε_i is independent of all other variables.

Now consider a regression of $\Delta\kappa_{1,i}$ on $[\kappa_{2,i}\hat{\pi}_i, 1] \otimes [1, v_i^{-2}]$, similar to what is in the text (the text includes more regressors, but they will be irrelevant here). Note that we have, combining what is above,

$$\Delta\kappa_{1,i} = a\kappa_{2,i}\hat{\pi}_i + bs_i\kappa_{2,i}\hat{\pi}_i + \varepsilon_i \quad (\text{A.17})$$

The population coefficient on $\kappa_{2,i}\hat{\pi}_i$ recovers a , the coefficient on $v_i^{-2}\kappa_{2,i}\hat{\pi}_i$ recovers b , and the estimates combined yield σ_i^{-2} (and hence also its cross-sectional mean). That is, the control function regression is correctly specified.

We can also consider a regression of $\Delta\kappa_{1,i}^2$. We have

$$\mathbb{E}[\Delta\kappa_{1,i}^2 | \kappa_{2,i}, v_i] = \mathbb{E}[\sigma_i^{-4}\kappa_{2,i}^2\hat{\pi}_i^2 + \varepsilon_i^2 + 2\varepsilon_i\sigma_i^{-2}\kappa_{2,i}\hat{\pi}_i | \kappa_{2,i}, v_i] \quad (\text{A.18})$$

$$= \mathbb{E}[\sigma_i^{-4}\kappa_{2,i}^2 | \kappa_{2,i}, v_i] \text{var}(\hat{\pi}_i) + \text{var}(\varepsilon_i) \quad (\text{A.19})$$

$$= \kappa_{2,i}^2 (a^2 + b^2s_i^2 + 2abs_i) \text{var}(\hat{\pi}_i) + \text{var}(\varepsilon_i) \quad (\text{A.20})$$

So if the regressors are $[\kappa_{2,i}^2, 1] \otimes [1, v_i^{-2}, v_i^{-4}]$, the regression is correctly specified and we can recover $E[\sigma_i^{-4}]$, which is the mean response of expected squared changes in beliefs to changes in uncertainty.

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

| (a) First stage, 1 year inflation | | |
|-----------------------------------|-----------------------------|-----------------------------|
| Dep var: | (1) Unc. \times Surpr. | (2) Unc. \times Surpr. |
| Tenure \times Surpr. | -0.325*** (-6.17) | -0.183*** (-3.91) |
| Infl. Surpr. | 13.817*** (18.49) | 12.228*** (24.32) |
| 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 | | |
|-----------------------------------|-----------------------------|-----------------------------|
| Dep var: | (1) Unc. \times Surpr. | (2) Unc. \times Surpr. |
| Tenure \times Surpr. | -0.266*** (-4.34) | -0.226*** (-4.48) |
| Infl. Surpr. | 13.748*** (19.14) | 12.631*** (28.32) |
| 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. Robust t -statistics allowing for two-way clustering by respondent and time are reported in parentheses. 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: | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ |
| Tenure \times Infl. Surpr. | -0.0061*** (-8.89) | -0.0025*** (-3.36) | -0.0056*** (-9.34) | -0.0027*** (-4.23) | -0.0064*** (-9.73) | -0.0031*** (-4.12) |
| Infl. Surpr. | 0.1510*** (16.86) | 0.0488*** (7.64) | 0.1419*** (18.91) | 0.0555*** (8.70) | 0.1650*** (23.90) | 0.0570*** (8.93) |
| 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: | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ |
| Tenure \times Infl. Surpr. | -0.0039*** (-5.35) | -0.0021*** (-2.70) | -0.0031*** (-4.71) | -0.0014** (-2.00) | -0.0040*** (-5.27) | -0.0024*** (-2.97) |
| Infl. Surpr. | 0.0887*** (11.73) | 0.0383*** (4.71) | 0.0800*** (11.50) | 0.0356*** (5.41) | 0.0964*** (14.00) | 0.0441*** (5.83) |
| 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 2. 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. Robust t -statistics allowing for clustering by respondent and time are reported in parentheses. 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: | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ |
| Tenure \times Infl. Surpr. | -0.0053*** (-6.89) | -0.0021*** (-2.94) | -0.0049*** (-6.94) | -0.0023*** (-3.50) | -0.0059*** (-7.36) | -0.0027*** (-3.48) |
| Infl. Surpr. | 0.1456*** (15.42) | 0.0456*** (7.44) | 0.1373*** (16.66) | 0.0527*** (7.87) | 0.1617*** (20.56) | 0.0538*** (8.50) |
| Tenure | 0.0067*** (3.20) | 0.0051*** (3.37) | 0.0055** (2.55) | 0.0038** (2.24) | 0.0033 (1.55) | 0.0035** (2.27) |
| 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: | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ | $\Delta\kappa_1$ |
| Tenure \times Infl. Surpr. | -0.0033*** (-4.06) | -0.0020** (-2.46) | -0.0024*** (-3.06) | -0.0012 (-1.64) | -0.0033*** (-3.67) | -0.0023*** (-2.73) |
| Infl. Surpr. | 0.0838*** (10.67) | 0.0372*** (4.49) | 0.0745*** (9.97) | 0.0341*** (5.01) | 0.0918*** (11.90) | 0.0435*** (5.52) |
| Tenure | 0.0062*** (3.18) | 0.0019 (1.10) | 0.0065*** (3.11) | 0.0020 (1.13) | 0.0046** (2.16) | 0.0007 (0.41) |
| 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.

Table A.4: Second Stage Regression: Coefficients

| | 1 year inflation | | 3 year inflation | |
|---|------------------------|-----------------------|------------------------|------------------------|
| | (1) | (2) | (3) | (4) |
| Unc. \times Infl. Surpr. | 0.0230*** (6.84) | 0.0114*** (3.44) | 0.0120*** (3.72) | 0.0032 (1.23) |
| Unc. $\times v_{i,t} \times$ Infl. Surpr. | 0.0000160*** (4.49) | 0.0000075** (2.29) | 0.0000166*** (5.75) | 0.0000110*** (3.09) |
| Unc. \times Infl. Surpr. ² | 0.0000411 (1.60) | -0.0000025 (-0.12) | 0.0000212 (1.08) | -0.0000076 (-0.36) |
| Infl. Surpr. | -0.0804*** (-3.60) | -0.0667** (-2.39) | -0.0458** (-1.99) | -0.0105 (-0.46) |
| Infl. Surpr. ² | -0.0017 (-1.18) | 0.0007 (0.82) | -0.0011 (-1.08) | 0.0004 (0.61) |
| $v_{i,t} \times$ Infl. Surpr. | -0.0247*** (-7.32) | -0.0128*** (-3.90) | -0.0135*** (-4.15) | -0.0047* (-1.76) |
| $v_{i,t}$ | 0.0069*** (6.98) | 0.0033*** (5.51) | 0.0057*** (7.84) | 0.0031*** (5.98) |
| Constant | 0.0921*** (4.08) | 0.0110 (0.94) | 0.0424*** (2.94) | 0.0028 (0.27) |
| 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 3.

Table A.5: Second Stage Regression: Robustness

(a) 1-Year Inflation

| | Dep. var.: $\Delta\kappa_{1,i,t}$ | | | | | |
|--|-----------------------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Unc. \times Infl. Surpr., avg marginal effect | 0.0287*** (6.81) | 0.0135*** (3.11) | 0.0266*** (7.14) | 0.0144*** (3.71) | 0.0294*** (7.05) | 0.0158*** (3.61) |
| $v_{i,t} \times$ Infl. Surpr., avg marginal effect | -0.0317*** (-7.48) | -0.0157*** (-3.65) | -0.0289*** (-7.74) | -0.0164*** (-4.28) | -0.0327*** (-7.80) | -0.0184*** (-4.27) |
| 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.: $\Delta\kappa_{1,i,t}$ | | | | | |
|--|-----------------------------------|-----------------------|-----------------------|----------------------|-----------------------|-----------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Unc. \times Infl. Surpr., avg marginal effect | 0.0183*** (4.43) | 0.0092** (2.53) | 0.0145*** (3.95) | 0.0059* (1.83) | 0.0181*** (4.34) | 0.0099*** (2.62) |
| $v_{i,t} \times$ Infl. Surpr., avg marginal effect | -0.0205*** (-4.95) | -0.0115*** (-3.11) | -0.0164*** (-4.46) | -0.0081** (-2.48) | -0.0205*** (-4.91) | -0.0126*** (-3.34) |
| 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 3, 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.

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

| | 1 year inflation | | 3 year inflation | |
|--|----------------------|-----------------------|----------------------|-----------------------|
| | (1) | (2) | (3) | (4) |
| Unc. \times Infl. Surpr., avg marginal effect | 0.0516** (2.26) | 0.0280*** (5.39) | 0.0318* (1.94) | 0.0374*** (5.12) |
| $v_{i,t} \times$ Infl. Surpr., avg marginal effect | -0.0523** (-2.28) | -0.0308*** (-5.75) | -0.0337** (-2.06) | -0.0400*** (-5.34) |
| 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 3, but using the respondents' point estimate of inflation to construct the dependent variable.

Table A.7: Second Stage Regression: Additional Controls

| | 1 year inflation | | 3 year inflation | |
|--|-----------------------|--------------------|---------------------|--------------------|
| | (1) | (2) | (3) | (4) |
| Unc. \times Infl. Surpr., avg marginal effect | 0.0157*** (4.50) | 0.0064** (1.97) | 0.0058* (1.80) | -0.0030 (-1.09) |
| $v_{i,t} \times$ Infl. Surpr., avg marginal effect | -0.0314*** (-4.00) | -0.0154 (-1.49) | -0.0148* (-1.87) | -0.0008 (-0.09) |
| 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 3, but with a larger set of controls: the Kronecker product of (inflation surprise; uncertainty \times inflation surprise) with (a constant, the control function component ($v_{i,t}$), the inflation surprise, and $v_{i,t}^{inv}$), where $v_{i,t}^{inv}$ is the inverse of ($v_{i,t} - \min(v_{i,t}) + 10$).

Table A.8: Second Stage Regression: 2SLS Estimates (Robustness Across Inflation Measures)

(a) 1-Year Inflation

| | Dep. var.: $\Delta\kappa_{1,i,t}$ | | | | | | | |
|----------------------------|-----------------------------------|----------------------|----------------------|--------------------|----------------------|----------------------|----------------------|--------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Unc. \times Infl. Surpr. | 0.0143*** (4.96) | 0.0116*** (2.76) | 0.0214*** (3.58) | 0.0160* (1.83) | 0.0199*** (3.97) | 0.0155*** (2.59) | 0.0282*** (2.85) | 0.0206* (1.89) |
| Infl. Surpr. | -0.0866*** (-2.73) | -0.0975** (-2.07) | -0.1874** (-2.28) | -0.1661 (-1.50) | -0.1607** (-2.43) | -0.1409** (-1.99) | -0.3016** (-2.02) | -0.2212 (-1.58) |
| 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.: $\Delta\kappa_{1,i,t}$ | | | | | | | |
|----------------------------|-----------------------------------|--------------------|---------------------|--------------------|--------------------|--------------------|--------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Unc. \times Infl. Surpr. | 0.0094*** (2.81) | 0.0032 (1.14) | 0.0165** (2.38) | 0.0084** (2.04) | 0.0142** (2.33) | 0.0065 (1.64) | 0.0216* (1.88) | 0.0105** (2.17) |
| Infl. Surpr. | -0.0687* (-1.74) | -0.0166 (-0.54) | -0.1766* (-1.77) | -0.0785 (-1.60) | -0.1363 (-1.63) | -0.0494 (-1.09) | -0.2678 (-1.48) | -0.1019* (-1.72) |
| 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 $\Delta\kappa_{1,i,t}$ (measured based on the bin mean) on the interaction of uncertainty and inflation surprises, $\text{Unc.} \times \hat{\pi}_{i,t}$, and on the inflation surprise $\hat{\pi}_{i,t}$. The endogenous regressor $\text{Unc.} \times \hat{\pi}_{i,t}$ is instrumented using tenure interacted with the inflation surprise, i.e., $\text{Tenure}_{i,t} \times \hat{\pi}_{i,t}$, 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. Standard errors are two-way clustered by respondent and time; t -statistics are reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.