Inflation Surprises and Perception of Inflation Risks

Miguel Acosta [*]
University of Wisconsin-Madison

Yeji Sung [†]
Federal Reserve Bank of San Francisco

March 20, 2025

Please do not cite or circulate without authors' permission.
Click for the latest version

Abstract

We document a robust relationship between inflation surprises and increases in perceived inflation uncertainty–even among professional forecasters and in stable, low-inflation environments. This pattern contradicts standard models of expectation formation, which predict that perceived uncertainty should remain stable when inflation itself is stable. We propose a new explanation: limited memory in the learning process. In our model, imprecise memory causes inflation surprises to have lasting effects on perceived uncertainty, even after extended periods of learning. This mechanism implies that investors demand higher compensation for holding assets with inflation-sensitive payoffs, shaping how perceived inflation risk is priced in financial markets.

[*]Department of Economics, University of Wisconsin-Madison. miguel.acosta@wisc.edu.
[†]Economic Research, Federal Reserve Bank of San Francisco. yeji.sung@sf.frb.org.

1 Introduction

Perception of risk is fundamental to numerous economic decision-making processes. When households and firms believe extreme events are more likely, they tend to postpone investment in durable goods. In financial markets, participants assess the risk associated with underlying cash flows, which is particularly evident in the pricing of financial contracts such as variance swaps and inflation swaps. Understanding the formation of risk perception is therefore crucial.

This paper examines how people's perception of risk deviates from optimal Bayesian inference. We suppose that individuals are subject to memory friction, causing risk perception to be based on imperfect recollections of past data rather than all available information. Our model predicts that perceived risk fluctuates perpetually, even when the true level of risk (or volatility of the process) remains constant. Moreover, following large variations in the underlying process, people tend to view extreme realizations as more likely in the future. We validate this model prediction using inflation forecast data from professional forecasters.

To illustrate this concept, we consider a case where the variance of a random variable following a Gaussian distribution is unknown, while its mean is known. To clarify the intuition, we suppose that both the mean and variance are constant. In this scenario, the history of deviations from the mean indicates the dispersion of the random variable. For a decision-maker (DM) learning about the variance by observing realized variations, the magnitude of surprises is informative in gauging the variance, regardless of whether they are recent or distant.

We propose a learning model where DM's perception is based on noisy memory, i.e., a noisy recollection of past data. The essence of our model is to conceptualize a stock of memory as noisy summary statistics. We show that lower-dimensional statistics summarize past data informative about the underlying variance. We then posit that this statistic is only available with noise in subsequent periods. This formulation nests the perfect memory benchmark, with the extent of the noise characterized by a single parameter. While we term this a model of noisy memory, it does not necessarily imply forgetfulness but rather formalizes the friction in accurately utilizing past data, which could capture various imperfections in information processing.

A key finding of our model is that DM overweights recent surprises compared to the perfect memory benchmark. Unlike a statistician who would assign equal weight to each realized variation when inferring variance, DM puts disproportionately high weight on recent experiences. Furthermore, weights on older data are stochastic and decay exponentially on average. This means the sequence of observation matters in shaping perceptions.

Importantly, such over-extrapolation persists even in the long run. Unlike learning with perfect memory, where new information eventually adds negligible value, out model shows that DM remains perpetually uncertain about the exact value of the variance due to noisy accumulation of knowledge. Consequently, DM continues to learn and adjust perceptions based on recent experiences even in the long run.

We test our model's main prediction using survey data of US professional forecasters. Our findings support the theory: when forecasters are surprised by recent inflation data, they become more

uncertain about future inflation. This positive relationship between the size of inflation surprises and perceived inflation risk holds not only during periods of substantial volatility—such the late 1970s—but also during less-volatile periods.[1]

Our proposed model offers an alternative explanation for fluctuating risk perception compared to the hypothesis of stochastic variance. While learning about a stochastically-varying variance can produce similar patterns between surprises and perceived risk, our model explains fluctuations in risk perception even when the volatility of the variance is negligible, or in our case, zero. This addresses a common challenge in the literature, where it has been difficult to justify variations in risk perception with realistic degrees of stochastic movement in the variance process (e.g., see Bakshi and Skoulakis (2010)).

2 Model

Environment A decision-maker (DM) forecasts a future realization of a variable y_t that follows a normal distribution with mean μ and precision ω:

$$y_t \sim \mathcal{N}\left(\mu,\, \omega^{-1}\right). \tag{1}$$

While the DM knows this variable is drawn from a normal distribution, the exact data-generating process—parameterized by μ and ω—is unknown. Our primary interest in this paper is on the DM's perception of the volatility of y_t, which we refer to as perceived volatility. To focus on how beliefs about $\frac{1}{\omega}$ are formed, we start with a more straightforward setup where the DM has a correct understanding of the mean μ. Thus, the DM needs to learn about one parameter: ω.

We assume that the following Gamma distribution describes the DM's prior belief about ω:

$$\omega \sim \Gamma\left(\alpha,\, \beta\right).$$

This prior belief ensures that the DM's belief about ω inferred from the history of y_t again follows a Gamma distribution. We define perceived volatility as $E_t\left[\omega^{-1}\right]$, where the expectation is conditional on information available to DM at time t as detailed below.

Available information The DM is Bayesian and thus forms their beliefs about the model's parameters using a combination of their prior information and any new information available to them. The DM learns about the model parameters by observing the realizations of y_t. We make the conventional assumption that at each period, y_t is realized, and the DM observes this realization without error.

Where this paper differs from conventional Bayesian learning is in the use of prior information. This paper studies how beliefs are formed when they are *not* accurately based on all past y_t: The

[1] While limitations in the survey's design do not allow us to study the most-recent inflation during COVID periods, our findings and model suggest that it was likely followed by an increase in the perceived inflation risk.

DM's beliefs about the parameters are not described as $\omega|\, y_1, \cdots, y_t$. Instead, past data is less accurately incorporated into the DM's current belief. To clarify this, we describe that a statistician's views are based on "perfect memory," while the DM's views are based on "noisy memory." In the following section, we discuss how noisy memory is modeled.

2.1 Learning with Perfect Memory

In this section, we consider how a statistician's views are formed. This will serve as the full information limit of our model. Suppose the $\mathcal{Y}^t$ denotes the history of y_t, that is, $\mathcal{Y}^t \equiv \{y_1, \cdots, y_t\}$. A statistician accurately infers the value of ω from the history of $\mathcal{Y}^t$.

It is well understood that the conditional distribution $\omega|\, \mathcal{Y}^t$ follows a Gamma distribution and that the distribution evolves in a recursive formula. Here, we summarize this point and introduce a notation. At any time t, we can describe the *prior* belief about ω at time t as

$$\omega|\, \mathcal{Y}^{t-1} \sim \Gamma\left(\alpha + \alpha^*_{t-1},\, \beta + \beta^*_{t-1}\right).$$

Then, an observation of y_t leads DM to have the *posterior* belief about ω as follows.

$$\omega|\, \mathcal{Y}^t \sim \Gamma\left(\alpha + \alpha^*_t + \frac{1}{2},\, \beta + \beta^*_t\right)$$

where α^*_t and β^*_t recursively evolve according to

$$\alpha^*_t = \alpha^*_{t-1} + \frac{1}{2} \tag{2a}$$

$$\beta^*_t = \beta^*_{t-1} + \frac{1}{2}\left(y_t - \mu\right)^2. \tag{2b}$$

Thus, the distribution $\omega|\, \mathcal{Y}^t$ is completely summarized by the two variables α^*_t and β^*_t. We can see that α^*_t increases incrementally independent of the realization of y_t, while β^*_t reflects how far off the realized y_t is from its expected value μ. Furthermore, the sequence of $\{\alpha^*_\tau, \beta^*_\tau\}^t_{\tau=0}$ is described below.

$$\alpha^*_t = \frac{t}{2} \tag{3a}$$

$$\beta^*_t = \frac{1}{2}\sum_{\tau=1}^{t}\left(y_\tau - \mu\right)^2 \tag{3b}$$

starting from the initial conditions $\alpha_0 = 0$ and $\beta_0 = 0$. Thus, α^*_t stores the length of the learning experiences, while β^*_t stores the stock of experienced surprises until time t.

The DM's perceived volatility at time t is given by $E[\omega^{-1}\big|\, \mathcal{Y}^t] = \frac{\alpha+\alpha^*_t}{\beta+\beta^*}$. Thus,

$$E[\omega^{-1}\big|\, \mathcal{Y}^t] = \frac{\beta + \frac{1}{2}\sum_{\tau=1}^{t}\left(y_\tau - \mu\right)^2}{\alpha + \frac{t}{2} - 1} \tag{4}$$

from equations (3). We define the weight on the variation realized in time τ as

$$s_\tau^* = \frac{\frac{1}{2}}{\alpha + \frac{t}{2} - 1}, \quad \tau \leq t. \tag{5}$$

As expected, we observe that all previously realized variation receives the same weight. That is, $s_\tau^* = s_t^*$ for all $\tau \leq t$. This means that all experiences have the same influence on DM's current perception of risk, regardless of how recent the experience is.

We can furthermore see that $s_t^* \to 0$ as $t \to \infty$. That is, recent observations will influence DM's perception of ω^{-1} less and less over time. By the law of large numbers, it is also straightforward to see that the estimate (4) converges to a true ω^{-1} as learning opportunities accumulate.

2.2 Learning with Noisy Memory

This section considers how the DM's views are formed based on noisy memory. In this case, the DM's beliefs about ω are not accurately based on the history of y_t. Thus, the DM's beliefs less accurately incorporate information from past data than a statistician's beliefs.

We propose a succinct and intuitive theory to introduce memory frictions. Following the work of Azeredo da Silveira et al. (2020) and Sung (2022), our model of memory will posit that the DM's knowledge about the variance of y_t is distilled down to a lower-dimensional set of sufficient statistics. This is motivated by the notion that the DM's beliefs are characterized by parametric distributions (here, the Gamma distribution), then those beliefs are completely summarized by the distributions parameters (in the case of the Gamma distribution, this means just two variables). Thus, these parameters summarize the knowledge about the history of past data realizations relevant for inferring the level of ω without losing any information.

For example, from the previous section, α_t^* and β_t^* completely describe the posterior distribution. Furthermore, β_t^* captures information from past realizations, as the other parameter α_t^* only depends on the number of observations. Thus, the inference based on β_t^* coincides with the statistician's view: It is sufficient for the DM to base her forecasts on β_t^* to replicate the perfect memory views of a statistician. Lemma 1 summarizes this point.

Lemma 1. *Suppose a posterior belief about ω is given as $\omega|\, \mathcal{Y}^t \sim \Gamma(\alpha + \alpha_t^*,\, \beta + \beta_t^*)$. Then, it must be that $\omega|\, \beta_t \sim \Gamma(\alpha + \alpha_t^*,\, \beta + \beta_t^*)$.*

Proof. The lemma follows from $\beta_t^*|\, \omega \sim \Gamma(\alpha_t^*,\, \omega)$. □

Inspired by the fact that β_t^* fully captures knowledge about ω in the perfect memory case, we suppose that the DM's views are based on noisy memory about this sufficient statistic. Thus, we propose that the DM's knowledge about ω until the beginning of time t is stored in a one-dimensional "memory variable" m_{t+1}, which evolves according to the following definition.

Definition 1. *If the time-t posterior belief about ω is given as $\omega|\, m_t, y_t \sim \Gamma(\alpha + \alpha_t,\, \beta + \beta_t)$, then*

memory variable m_{t+1} is defined as

$$m_{t+1} = \beta_t \, u_{t+1} \tag{6}$$

where u_{t+1} is an i.i.d. draw from the Beta distribution $\mathcal{B}\left(\lambda\,\alpha_t,\ (1-\lambda)\,\alpha_t\right)$ for a given scalar $\lambda \in (0,1)$.

A single parameter λ describes the extent of noisy memory. We can see this from the noise term u_{t+1}. Note that $E[u_{t+1}] = \lambda$ and $V[u_{t+1}] = \frac{\lambda(1-\lambda)}{\alpha+1}$. Thus, the perfect memory case is nested as $\lambda \to 1$ since $m_{t+1} = \beta_t$, in which case the DM's views coincide with a statistician who accurately uses all available past y_t. On the other hand, the formulation above also incorporates the case in which the memory variable m_{t+1} carries no information about ω. When $\lambda \to 0$, we can see that $m_{t+1} = 0$. Thus, m_{t+1} contains no information about previous realizations that are informative about the variability of y_t. In other words, we have that $\omega|\, m_{t+1} \sim \Gamma\left(\alpha,\, \beta\right)$ as m_{t+1} does not improve DM's initial prior about ω.

Given equation (6), the DM's beliefs about ω evolve as described in Proposition 1. Importantly, the proposed memory structure allows the prior belief about ω to be the conjugate prior, as was the case in the perfect memory benchmark. Thus, the DM's prior and posterior beliefs about ω are described as Gamma distributions.

Proposition 1. *If memory is formed according to (6), then the time-t prior belief can be expressed as*

$$\omega|\, m_t \sim \Gamma\left(\alpha + \alpha_t^m,\, \beta + \beta_t^m\right)$$

The posterior belief, after observing y_t, evolves to

$$\omega|\, m_t, y_t \sim \Gamma\left(\alpha + \alpha_t,\, \beta + \beta_t\right)$$

where $\alpha_t = \alpha_t^m + \frac{1}{2}$ and $\beta_t = \beta_t^m + \frac{1}{2}\left(y_t - \mu\right)^2$. Finally, the beginning-of-period prior at $t+1$ is derived as

$$\omega|\, m_{t+1} \sim \Gamma\left(\alpha + \alpha_{t+1}^m,\, \beta + \beta_{t+1}^m\right)$$

where $\alpha_{t+1}^m = \lambda\,\alpha_t$ and $\beta_{t+1}^m = m_{t+1} = \beta_t\, u_{t+1}$.

Proof. See Appendix A. □

As Proposition 1 suggests, DM's posterior belief about ω follows a Gamma distribution specified by two variables, α_t and β_t. They evolve recursively as follows.

$$\alpha_t = \lambda\,\alpha_{t-1} + \frac{1}{2} \tag{7a}$$

$$\beta_t = \beta_{t-1}\, u_t + \frac{1}{2}\left(y_t - \mu\right)^2 \tag{7b}$$

where u_t is randomly drawn from $\mathcal{B}(\lambda\,\alpha_{t-1},\,(1-\lambda)\,\alpha_{t-1})$. Compared to the perfect memory counterpart (2), there are important differences in how α_t and β_t evolve. First, while the evolution of α_t from α_{t-1} is still independent of the realization of y_t, the change is somewhat attenuated since $\lambda \in (0,1)$. Second, while the fluctuation in β_t is still tied to the realized variation $(y_t-\mu)^2$, there is a second source of fluctuation: The stochastic term u_t is independent of y_t and multiplies β_{t-1}. Thus, the experience of past surprises summarized in β_{t-1} influences β_t in a stochastic manner.

By recursively iterating (7), we can describe α_t and β_t as follows.

$$\alpha_t = \frac{1}{2} + \frac{\lambda}{2} + \cdots + \frac{\lambda^{t-1}}{2} = \frac{1}{2}\,\frac{1-\lambda^t}{1-\lambda} \tag{8a}$$

$$\beta_t = \frac{1}{2}\sum_{\tau=1}^{t} (y_\tau-\mu)^2\;\bar{u}_{t,\tau} \tag{8b}$$

starting from the initial points $\alpha_0 = \beta_0 = 0$. [2] The term $\bar{u}_{t,\tau}$ is defined as $\Pi_{i=\tau+1}^{t} u_i$ for $\tau \leq t-1$ and $\bar{u}_{t,t} = 1$. The equations (8) summarize how limited memory affects the DM's learning. Note first that when $\lambda = 1$, α_t and β_t converge to their perfect-memory counterpart α_t^* and β_t^* described in (3). However, when memory is limited, α_t does not linearly increase with the length of learning opportunities; α_t converges to a positive scalar as $t \to \infty$. Also, β_t is not simply summing past realized variations. Previous surprises are multiplied by idiosyncratic noise, and the noise is cumulative, as captured by $\bar{u}_{t,\tau}$. Therefore, older surprises have a less precise impact on the DM's estimate of ω.

The DM's perceived variance is derived as

$$E[\omega^{-1}|\, m_t, y_t] = \frac{\beta + \frac{1}{2}\sum_{\tau=1}^{t}(y_\tau-\mu)^2\;\bar{u}_{t,\tau}}{\alpha + \frac{1}{2}\frac{1-\lambda^t}{1-\lambda} - 1} \tag{9}$$

from equations (8). As before, we define the weight on the variation realized at time τ as

$$s_\tau = \frac{\frac{1}{2}\,\bar{u}_{t,\tau}}{\alpha + \frac{1}{2}\frac{1-\lambda^t}{1-\lambda} - 1}, \quad \tau \leq t. \tag{10}$$

Note that the weights on the past realizations are stochastic due to the compounded idiosyncratic noise $\bar{u}_{t,\tau}$. Furthermore, we can see that $s_\tau = \bar{u}_{t,\tau}\,s_t$ where s_t is the weight on the most recent variation. Thus, it is straightforward to see that

$$E[s_\tau] = \lambda^{t-\tau}\,s_t, \quad \tau \leq t \tag{11}$$

where the expectation is taken over possible realizations of $\bar{u}_{t,\tau}$. This implies that weights on past variations exponentially decay on average.

We can also see that the weight on the most recent observation s_t does not converge to zero as $t \to \infty$. Instead, the limit converges to a positive number tied to the extent of noisy memory.

[2] At $t = 1$, we assume there is no memory to be inherited. The posterior is then $\omega|\, y_1$ which is a Gamma distribution.

This means that recent surprises keep affecting the DM's perception of ω even after infinitely long learning opportunities. An extreme case of "no memory" clarifies intuition. Suppose the DM's beliefs are formed only based on the current observations despite having observed past variations of y_t. In this case, the DM's beliefs will markedly differ from those of a statistician, although they have observed the same data. Since past data does not affect the DM's beliefs ($\lambda = 0$), it must be that $\alpha_t = \frac{1}{2}$ and $\beta_t = \frac{1}{2}(y_t - \mu)^2$ at any t. Therefore, we can see that $E[\omega^{-1} \big| m_t, y_t]$ and $V[\omega^{-1} \big| m_t, y_t]$ keep fluctuating even in the long run. Instead of beliefs converging to a number, they inherit the randomness of the realized y_t.

2.3 Illustrating the Implications of Noisy Memory

In this section, we discuss a numerical example to emphasize the model predictions. We consider the case where y_t is drawn from $\mathcal{N}(0, 1)$. We suppose that DM's initial prior belief about ω^{-1} is characterized by $E[\omega^{-1}] = 1$ and $V[\omega^{-1}] = 0.5$, which implies $\alpha = 4$ and $\beta = 3$.

Figure 1 shows that recent experiences have bigger influences on the DM's perceived volatility when subject to noisy memory. The top figure describes the average weight put on variations realized in the past when the DM has 20 periods of observations. The x axis is the lag since $t = 20$, so the number 0 and 20 correspond to the most recent and the oldest time periods. The perfect memory case is described by the purple line with a marker. As expected from (5), old and new variations all receive the same weights. In inferring the variance of y_t, what matters is how big the extent of surprises are, not when they occurred. In comparison, this prediction does not hold under noisy memory, as captured by the solid blue and dashed green lines. More-recent observations receive higher weights compared to the perfect memory case. Furthermore, weights on older observations decay exponentially, as described by (10). This implies that the sequence of observations starts to matter if the DM's perception is based on noisy memory. In the bottom figure, we compare the perceived volatility, $E[\omega^{-1} \big| m_t, y_t]$, under two scenarios. In the first scenario, we suppose that $y_\tau = \mu + \omega^{-\frac{1}{2}}$ for the first 10 periods, followed by $y_\tau = \mu$ in the remaining periods. In the second scenario, we reverse the order. We call these two sequences "Surprise Earlier" and "Surprise Later." Note that if perception is based on perfect memory, the two sequences result in the same level of risk perception at $t = 20$. This is because the order in which surprises are realized should not matter. In comparison, when perception is based on noisy memory, the DM perceives the level of volatility is higher when large variations are recently realized. This prediction is a direct consequence of the decaying weights shown in the top figure.

Figure 2 illustrates the DM's uncertainty about the variance of y_t. For both figures, we report the *expected* level of the posterior uncertainty, averaging over the possible values of $V[\omega^{-1} \big| m_t, y_t]$. As discussed earlier, the posterior uncertainty fluctuates in tandem with the realized variation in y_t. The top figure shows how the posterior uncertainty about ω^{-1} evolves from $t = 0$. Despite observing the same data realizations, the DM's uncertainty about ω^{-1} varies depending on the extent of noisy memory λ. The purple bottom line (labeled with $\lambda = 1.00$) is the perfect memory benchmark. As expected, as learning opportunities accumulate, the DM's uncertainty about ω^{-1}

decreases over time and eventually reaches zero. Thus, the DM can learn the actual value of ω through learning. The other lines depict the noisy memory case, with varying degrees of noise in memory. The top red line (labeled with $\lambda = 0.00$) is the no-memory case discussed earlier. Knowledge from past observations is completely lost between periods, so the DM is always as uncertain about ω^{-1} as they are in the initial learning period. We can see that learning is slower under "noisier" memory. The bottom figure makes this point as well. We report the average values of the posterior uncertainty after a long learning period for each level of noisy memory on the x-axis. Noisier memory leads to higher uncertainty about the variance of y_t.

Figure 3 displays the impulse response function. We report how the posterior mean of ω^{-1}—the DM's perceived volatility—responds *in the long run* to a positive one standard deviation shock. Different lines correspond to varying degrees of noisy memory λ. The purple line (with a circle marker) is the perfect memory benchmark. As discussed earlier, after a long enough learning period, the DM eventually learns about ω. This means that the DM's perception of ω^{-1} does not change in response to a fluctuation in the realized y_t. This prediction does not hold when $\lambda < 1$. For all the rest of the lines, we see that the DM perceives that y_t is more variable after observing an unusually high realization of y_t. The degree of on-impact response and the dynamics afterward vary depending on the noise in memory. For example, in the no-memory case (the red dotted line), we see the largest on-impact response that immediately dies out the following period. Since the DM is hugely uncertain about the value of ω, a large variation of y_t signals that the variance of y_t is high. This inference, however, is not carried over to the following periods since her future beliefs do not put weight on past data. The posterior mean is serially correlated in the intermediate case of $\lambda \in (0, 1)$. A one-time experience of high y_t keeps influencing the DM's perception of ω^{-1}.

2.4 Testable predictions

The main prediction of the model we would like to test is captured by the relationship between perceived risk and recent surprises.

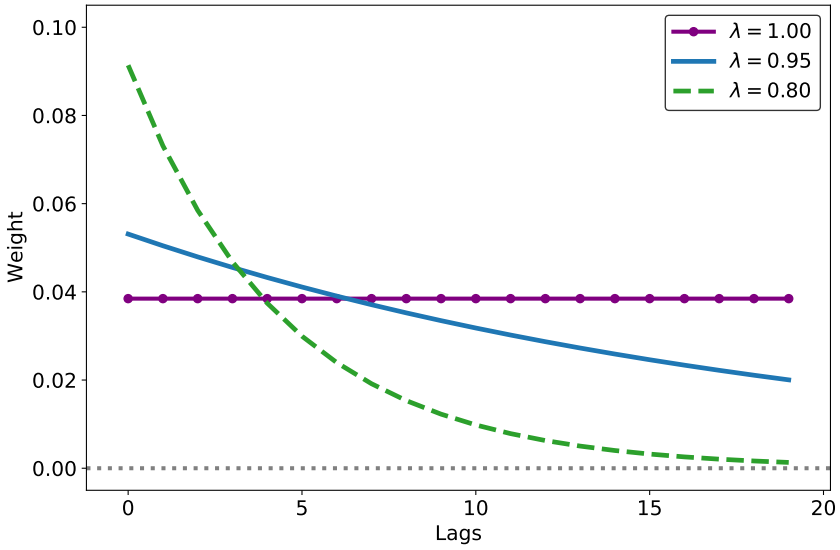
$$\beta \equiv \frac{Cov\left[PerceivedRisk_t, RecentSurprises_t\right]}{V[RecentSurprises_t]} \tag{12}$$

In our model, perceived volatility is defined as $E[\omega^{-1} \big| m_t, y_t]$, and recent surprises are defined as $(y_t - \mu)^2$. From (4) and (9), we can immediately see that $\beta = s_t$, where s_t is the weight on the most recent surprises as defined in (5) and (10).

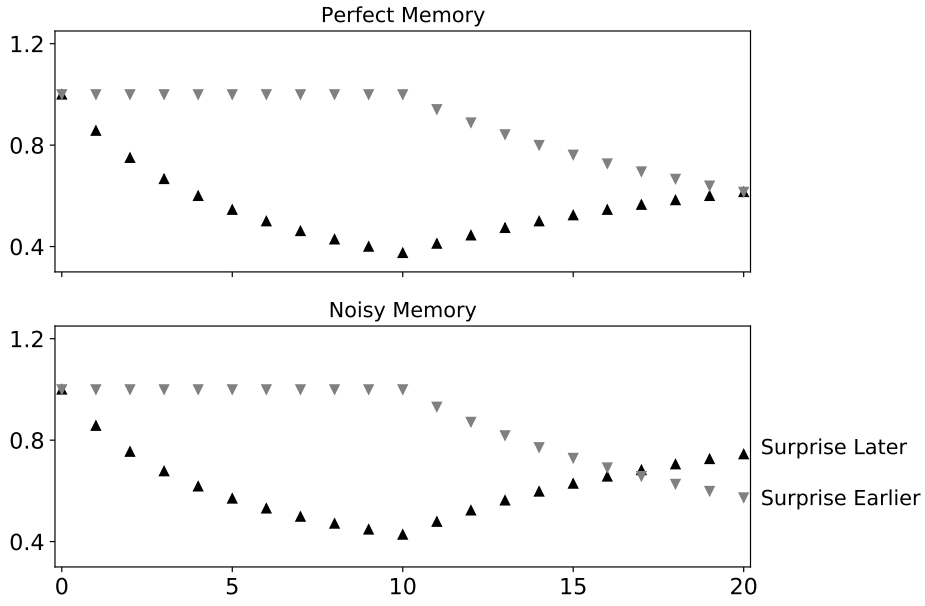
Importantly, the long-run properties of s_t allow us to test if our noisy memory theory is validated in the data. We have shown that β must be zero if DM is learning about the variance with perfect memory. In other words, perceived risk does not fluctuate with recent surprises. This is because the DM comes to have accurate knowledge of the actual variance. In contrast, under noisy memory, it

Figure 1: More Weight on Recent Surprises

(a) Weights on Past Surprises



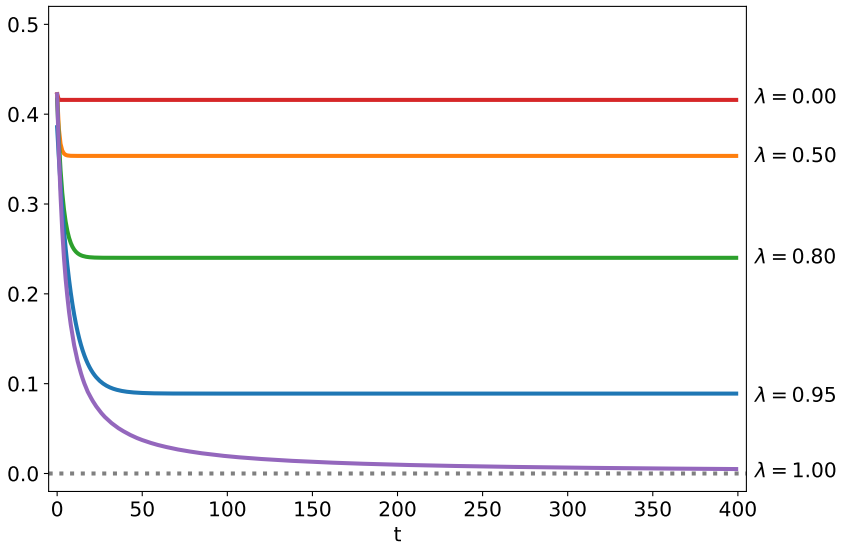
(b) Risk Perception and Order Effects



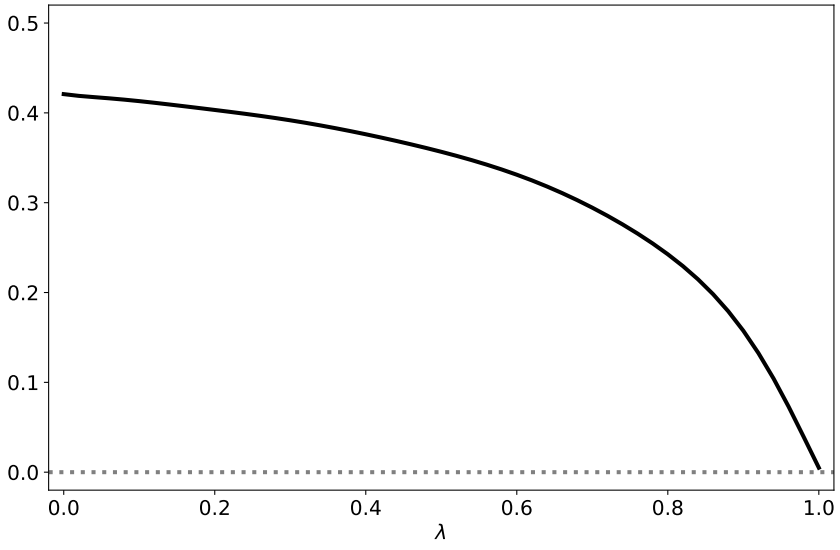
The top figure shows the weights on past surprises. The figure assumes that DM has 20 periods of learning opportunities. Each line corresponds to a different degree of noisy memory λ. The bottom figure reports $E[\omega^{-1}|\, m_t, y_t]$ when DM has observed the following sequences. The first scenario, labeled as "Surprise Earlier", assumes $y_\tau = \mu + \omega^{-\frac{1}{2}}$ for $\tau \leq 10$ and $y_\tau = \mu$ for $\tau > 20$. The second scenario, labeled as "Surprise Later", assumes $y_\tau = \mu$ for $\tau \leq 10$ and $y_\tau = \mu + \omega^{-\frac{1}{2}}$ for $\tau > 20$.

Figure 2: Posterior Uncertainty about ω^{-1}

(a) Evolution of $V\left[\omega^{-1}|m_t, y_t\right]$

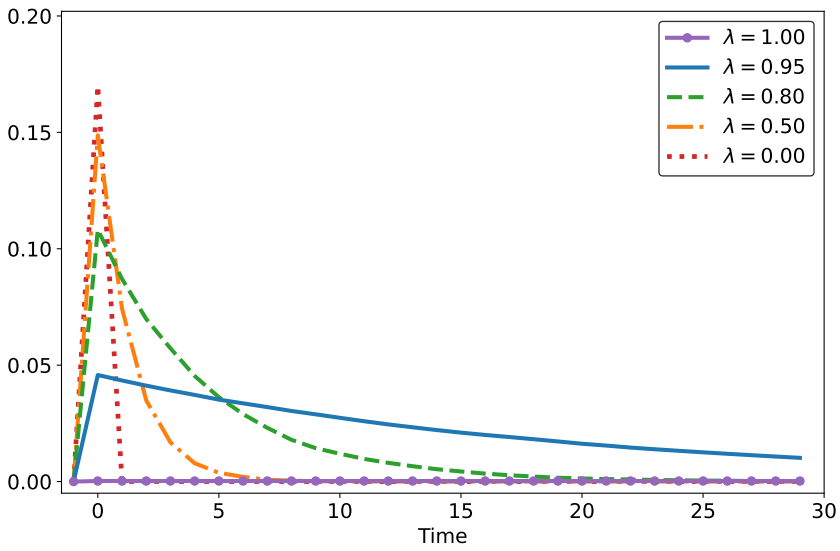


(b) Steady-state values of $V\left[\omega^{-1}|m_t, y_t\right]$



Above figures display the model-predicted posterior uncertainty of ω^{-1}. The top panel shows how the posterior uncertainty evolves for varying degrees of noisy memory λ. The bottom panel shows the steady-state levels of posterior uncertainty for each λ on the x axis.

Figure 3: Impulse Response Function of $E\left[\omega^{-1}|m_t, y_t\right]$



The figure displays how the posterior mean of ω^{-1} responds to a one standard deviation increase in the realization of y_t. Each line corresponds to a different degree of noisy memory λ. The figure assumes that DM has had infinite learning periods.

must be that

$$\beta = \left(2\left(\alpha - 1\right) + \frac{1}{1-\lambda}\right)^{-1} \tag{13}$$

in the long run. Thus, s_t is positive and more positive for noisier memory. With noisy memory, recent surprises continue influencing the DM's view about volatility. Thus, following a larger magnitude of surprises, the DM perceives that extreme realizations are more likely in the future. When the realized surprise is modest, the perceived volatility is lower under noisy memory compared to the perfect memory case. The slope of this relationship is steeper with noisier memory.

3 Empirical Evidence

We now turn to an empirical evaluation of the model's implications.

3.1 Data

Our data on inflation forecasts come from U.S. Survey of Professional Forecasters (SPF), administered quarterly by the Philadelphia Fed. For a subset of variables, the SPF asks forecasts to submit both point forecasts and a probability density function over realizations of those variables at particular horizons. Probability assessments are available for inflation based on the gross domestic product (GDP) deflator, the core personal-consumption expenditures price index, and the core consumer price index. The latter two series only begin in 2007, so we focus on expectations of inflation based on the GDP deflator, expectations of which have been collected since 1969.

We also collect data on real-time estimates of inflation from the Philadelphia Fed. The SPF is administered shortly after the "advance" release of GDP (including its componenets and price index). The release for quarter $t-1$ occurs late in the first month of quarter $t-1$. We denote the first-release of annualized inflation for quarter $t-1$ as $\pi_{t-1|t}$.[3] The Survey's administrators make this information available to forecasters when taking the survey at time t—it is, thus, the most recently available data on inflation.

Participants in quarter t submit a probability distribution over the level of inflation that will prevail at time τ. Specifically, participants are given B bins, indexed by b, that cover a finite range of the real number line. Each bin has a maximum u_b and minimum ℓ_b value.[4] The probability assigned by participant i that annual inflation at time τ will fall in bin b is given by $\pi_{i,t,\tau,b}$. We denote the midpoint of the bin by $m_b = 0.5 \times (u_b + \ell_b)$ and define

$$\mathbb{M}_{i,t,\tau} = \sum_{b=1}^{B} \pi_{i,t,\tau,b} m_b,$$

and their subjective inflaiton uncertainty by

$$\mathbb{V}_{i,t,\tau} = \sum_{b=1}^{B} \pi_{i,t,\tau,b} \left(m_b - \mathbb{M}_{i,t,\tau}\right)^2.$$

This is our primary measure of inflation uncertainty. In practice, τ is either the end of year t, or the following year. Our primary measure of the inflation surprise is given by

$$x_{i,t} \equiv \pi_{t-1|t} - \mathbb{E}_{t-1}\left[\pi_{t-1}\right],$$

[3] We construct this from real-time estimates of the prices index, P. Therefore, $\pi_{t-1|t} = \left(\left(\frac{P_{t-1|t}}{P_{t-2|t}}\right)^4 - 1\right) \times 100$. We use this same construction when working with forecasts of the price level.

[4] This range is top- and bottom-coded. For the bottom-coded and top-coded bins, we set $u_b = \ell_b$. Thus, our estimate of the variance is likely downward biased. We are exploring alternative approaches in ongoing work.

Table 1: Summary Statistics

Baseline Sample (1992–2019)

	Observations	Mean	Std. Dev.	25th Pctile.	75th Pctile.
$s_{i,t}$	6349	-0.18	1.16	-0.86	0.52
$s_{i,t}^2$	6349	1.38	3.14	0.12	1.42
$s_{i,t}^2$, Q1 or Q2	3149	1.43	2.83	0.11	1.48
$s_{i,t}^2$, Q3 or Q4	3200	1.33	3.41	0.13	1.38
$\mathbb{V}_{i,t,\tau}$	6349	0.51	0.55	0.20	0.65
$\mathbb{V}_{i,t,\tau}$, Q1 or Q2	3149	0.57	0.55	0.24	0.73
$\mathbb{V}_{i,t,\tau}$, Q3 or Q4	3200	0.46	0.54	0.16	0.56
$\mathbb{V}_{i,t,\tau}$, current year	3191	0.41	0.47	0.14	0.52
$\mathbb{V}_{i,t,\tau}$, next year	3158	0.61	0.60	0.25	0.77

Extended Sample Sample (1969–2019)

	Observations	Mean	Std. Dev.	25th Pctile.	75th Pctile.
$s_{i,t}$	9363	-0.09	1.52	-0.92	0.71
$s_{i,t}^2$	9363	2.31	8.27	0.16	2.05
$s_{i,t}^2$, Q1 or Q2	4709	2.44	6.65	0.15	2.36
$s_{i,t}^2$, Q3 or Q4	4654	2.18	9.63	0.16	1.97
$\mathbb{V}_{i,t,\tau}$	9363	0.65	0.75	0.21	0.81
$\mathbb{V}_{i,t,\tau}$, Q1 or Q2	4709	0.72	0.79	0.25	0.87
$\mathbb{V}_{i,t,\tau}$, Q3 or Q4	4654	0.58	0.69	0.19	0.75
$\mathbb{V}_{i,t,\tau}$, current year	5486	0.56	0.68	0.19	0.70
$\mathbb{V}_{i,t,\tau}$, next year	3877	0.77	0.81	0.26	0.95

In words, this is the surprise component of (annualized) quarterly inflation from quarter $t-2$ to quarter $t-1$. We construct the term $\mathbb{E}_{t-1}[\pi_{t-1}]$ from the point forecasts of the price level, as this is the point forecast solicited by the survey.

Our baseline sample runs from 1992 through 2019. We start in 1992 because the definition of the exact variable being forecast changed. Before 1992, the SPF asked participants about gross national product. Starting in 1992, participants were asked about GDP. We consider an earlier start date (the earliest possible—1969—in a robustness exercise). Starting the sample after the extreme volatility in inflation in the late 1970s/early 1980s also ensures that we are learning about the behavior of forecasts when inflation is relatively stable. We end our analysis before the COVID pandemic for a practical reason—the increase in inflation in 2021 caused many forecasters to put high probability mass in the highest bin. This mechanically makes uncertainty low, but is in fact the result of unchanging bins in the survey.

Table 1 shows basic summary statistics of the variables we have constructed. Surprises are close to zero on average, but with a fairly wide standard deviation and, thus, average absolute size. The average squared forecast revision in our baseline sample is about one and a half percentage points (in annualized inflation terms). Surprises and uncertainty tend to be larger in the earlier part of the sample, consistent with the very volatile underlying processes in the late 70s. Uncertainty deceases later in the calendar year in both samples—hence the quarter-of-year fixed effects—though this is not a significant feature of the size of forecast revisions, since these are not calendar-time forecasts. Uncertainty about the current calendar year is also lower than uncertainty about the following year.

3.2 Estimates

To assess the relationship between inflation surprises and perceived inflation risk, we estimate the following specification

$$\mathbb{V}_{i,t,\tau} = \beta x_{i,t}^2 + \delta_i + \delta_{q(t),\tau} + \varepsilon_{i,t}. \tag{14}$$

The term δ_i denotes a forecaster fixed effect, while $\delta_{q(t),\tau}$ is a forecast-horizon $\times$ quarter-of-year fixed effect. We include the latter in order to address the fact that SPF participants are forecasting calendar-time outcomes, so $\tau - t$ becomes smaller later in the quarter, and can thus reduce uncertainty mechanically. In robustness exercises, we show that only focusing on next-year forecasts yields similar results. The coefficient of interest is β—the coefficient on the squared value of the inflation surprise. We scale this term by its unconditional standard deviation in the 1969–2019 sample for ease of interpretation. The theory of section 2 suggests that $\beta > 0$.

Table 2 shows the results. In the first column, we show our baseline estimates during the 1992–2019 sample, using all forecast horizons. The relationship between inflation surprises and perceived inflation risk is positive and statistically significant. In terms of magnitudes, a 1 standard deviation squared inflation surprise leads to an increase in the perceived variance of inflation of 0.05 percentage point. This is roughly ten percent of the average value of the perceived variance.

The additional columns of table 2 confirm that these findings are robust to the alternative spec-

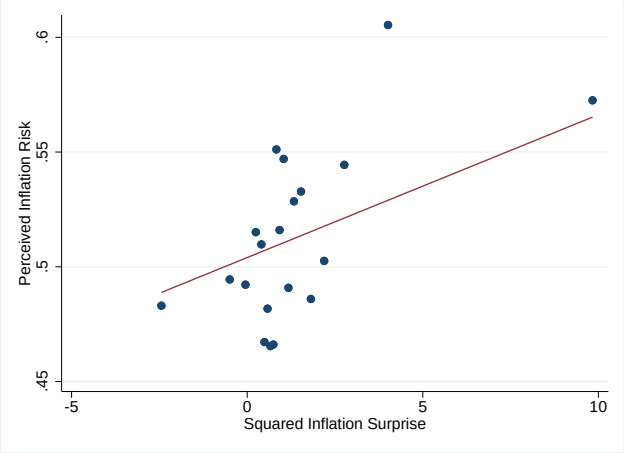
Table 2: Empirical Estimations: Inflation Surprises and Perceived Inflation Risk

s_{it}^2	0.0514	0.0431	0.0598	0.0322
	(0.0218)	(0.0204)	(0.0262)	(0.0110)
Observations	6349	3179	3145	9363
R^2	0.531	0.500	0.581	0.474
Sample	1992–2019	1992–2019	1992–2019	1969–2019
Horizons	0 & 1	0	1	0 & 1

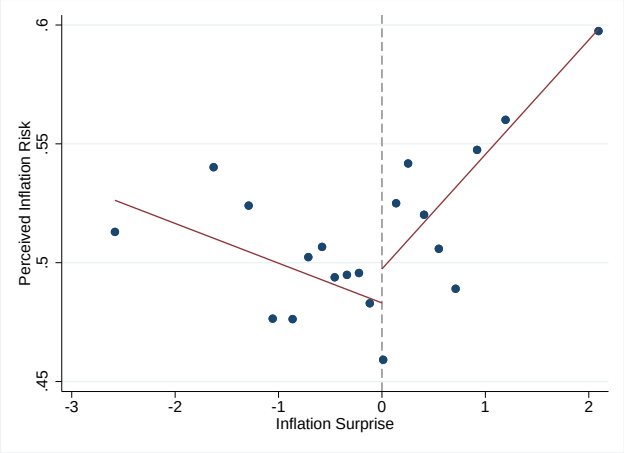
NOTE. This table shows estimates of equation (14). The left-hand side is expressed in percentage points, while $x_{i,t}^2$ has been scaled by its unconditional standard deviation (which can be seen in the second panel of table 1). The first column shows our baseline estimates. The second and third columns show our results for forecasts of inflation in the year the forecast is made, and in the subsequent year, respectively. The final column includes all observations since 1969. Standard errors clustered by i and t are shown in parentheses. All regressions include the fixed effects mentioned in the description of equation (14).

Figure 4: Scatter Plot of Baseline Regression

(a) Squared Surprises



(b) Surprises



NOTE. Panel (a) shows a binned scatter plot of equation(14) (i.e., all variables have been orthogonalized to the controls in that regression). Panel (b) allows positive and negative revisions to enter with different coefficients. See the note to table 2 for more details.

ifications considered. The second and third columns split the original sample by forecast horizon—either the current calendar year, or the year after. The results are attenuated somewhat for the current calendar year, though it is worth noting that perceived inflation risk in the current calendar year is much lower for the current calendar year, as seen in table 1. The results using the extended sample are similar to the baseline sample, though are also somewhat attenuated.

Figure 4 shows a scatter plot of the data underlying our baseline estimates. Both panels in the figure show inflation surprises and inflation risk orthogonalized to the same controls as in our baseline. Panel (a) shows squared surprises, as in our baseline estimates of equation (14).

Because squaring surprises can skew the importance of larger surprises (e.g. the point on the far right of panel (a)), in appendix B we show that this relationship continues to hold when using the absolute value of forecast revisions, and when estimating "Huber-robust" regressions, as in the work of Coibion et al. (2018). Panel (b) shows the non-squared surprises, and allows the coefficients on negative and positive values to be equal, in contrast to our baseline estimates. Despite this relaxation, the same pattern emerges—larger forecast revisions of both signs are associated with larger inflation uncertainty.

4 Conclusion

We propose a theory of how risk perception is formed when subject to noisy memory. We show that perceived risk fluctuates even when the underlying true risk does not change at all. Moreover, we show that recent surprises have an outsized influence on one's perceived risk, while older surprises have an exponentially decreasing influence.

Our theory provides a new implication for the asset pricing literature. Perceived risk has been argued to be important in explaining the historical equity premium and risk-free return. In this literature, the conventional approach is to consider a shock to the variance of the underlying process. For example, Weitzman (2007) argues that such variance shock can generate fluctuating market-wide risk perception, which explains several "puzzles" in the literature. While plausible, this idea did not have a quantitative success. As Bakshi and Skoulakis (2010) argues, explaining features of the asset prices with the variance shock still requires an implausible level of stochasticity that the data does not justify. Our model provides another reason why risk perception fluctuates, and this reason is not tied to the data generating process of the underlying variance.

References

AZEREDO DA SILVEIRA, R., Y. SUNG, AND M. WOODFORD (2020): "Optimally Imprecise Memory and Biased Forecasts," Working Paper 28075, National Bureau of Economic Research.

BAKSHI, G. AND G. SKOULAKIS (2010): "Do subjective expectations explain asset pricing puzzles?" *Journal of Financial Economics*, 98, 462–477.

BASSETTI, F., R. CASARIN, AND M. DEL NEGRO (2023): "Chapter 15 - Inference on probabilistic surveys in macroeconomics with an application to the evolution of uncertainty in the survey of professional forecasters during the COVID pandemic," in *Handbook of Economic Expectations*, ed. by R. Bachmann, G. Topa, and W. van der Klaauw, Academic Press, 443–476.

BEECHEY, M. J. AND J. H. WRIGHT (2009): "The high-frequency impact of news on long-term yields and forward rates: Is it real?" *Journal of Monetary Economics*, 56, 535–544.

BEN-DAVID, I., E. FERMAND, C. M. KUHNEN, AND G. LI (2018a): "Expectations Uncertainty and Household Economic Behavior," Working Paper 25336, National Bureau of Economic Research.

——— (2018b): "Expectations Uncertainty and Household Economic Behavior," Working Paper 25336, National Bureau of Economic Research.

BOLLERSLEV, T., G. TAUCHEN, AND H. ZHOU (2009): "Expected Stock Returns and Variance Risk Premia," *The Review of Financial Studies*, 22, 4463–4492.

COIBION, O., D. GEORGARAKOS, Y. GORODNICHENKO, G. KENNY, AND M. WEBER (2021): "The Effect of Macroeconomic Uncertainty on Household Spending," Working Paper 28625, National Bureau of Economic Research.

COIBION, O., Y. GORODNICHENKO, AND S. KUMAR (2018): "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review*, 108, 2671–2713.

COLLIN-DUFRESNE, P., M. JOHANNES, AND L. A. LOCHSTOER (2017): "Asset Pricing When 'This Time Is Different'," *Review of Financial Studies*, 30, 505–535.

DEW-BECKER, I., S. GIGLIO, A. LE, AND M. RODRIGUEZ (2017): "The price of variance risk," *Journal of Financial Economics*, 123, 225–250.

EVANS, G. AND S. HONKAPOHJA (2011): "Learning as a Rational Foundation for Macroeconomics and Finance," CEPR Discussion Papers 8340, C.E.P.R. Discussion Papers.

KOZLOWSKI, J., L. VELDKAMP, AND V. VENKATESWARAN (2020): "The Tail That Wags the Economy: Beliefs and Persistent Stagnation," *Journal of Political Economy*, 128, 2839–2879.

LOCHSTOER, L. A. AND T. MUIR (2022): "Volatility Expectations and Returns," *The Journal of Finance*, 77, 1055–1096.

STELLA, A. AND J. H. STOCK (2015): "14A state-dependent model for inflation forecasting," in *Unobserved Components and Time Series Econometrics*, Oxford University Press.

SUNG, Y. (2022): "Macroeconomic Expectations and Cognitive Noise," Working Paper.

WEITZMAN, M. L. (2007): "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, 97, 1102–1130.

A Detailed Derivations

We derive how beliefs about ω (the inverse of the variability of y_t) evolve if the "memory process" is formed according to (6). We show this recursively.

At $t = 1$, DM starts with a prior belief denoted as $\omega \sim \Gamma(\alpha,\, \beta)$. After observing y_1, the posterior belief is formed as

$$\omega|\, y_1 \sim \Gamma\left(\alpha + \underbrace{\frac{1}{2}}_{\equiv\alpha_1},\, \beta + \underbrace{\frac{1}{2}(y_1 - \mu)^2}_{\equiv\beta_1}\right)$$

from which we can see that $\beta_1|\,\omega \sim \Gamma(\alpha_1,\, \omega)$. The memory variable is defined as

$$m_2 = \beta_1\, u_2$$

where $u_2 \sim \mathcal{B}(\lambda\alpha_1,\, (1-\lambda)\,\alpha_1)$. Then, it must be that the likelihood function of m_2 is derived as $m_2|\,\omega \sim \Gamma(\lambda\,\alpha_1,\, \omega)$.

Given this structure, we can derive the prior belief at $t = 2$ as follows.

$$\begin{aligned}\omega|\, m_2 &\propto f(m_2|\,\omega)\; f(\omega) \propto \left(\omega^{\lambda\,\alpha_1}\, \exp^{-\omega\, m_2}\right)\left(\omega^{\alpha-1}\, \exp^{-\beta\,\omega}\right)\\ &\propto \Gamma\left(\alpha + \underbrace{\lambda\,\alpha_1}_{\equiv\alpha_2^m},\, \beta + \underbrace{m_2}_{=\beta_1\, u_2 \equiv \beta_2^m}\right)\end{aligned}$$

Thus, the memory variable m_2 maintains the prior distribution to be a Gamma distribution. I use the notation $\omega|\, m_2 \sim \Gamma(\alpha + \alpha_2^m,\, \beta + \beta_2^m)$ to denote this distribution. After observing y_2, DM updates her beliefs, which again are described as a Gamma distribution.

$$\omega|\, m_2, y_2 \sim \Gamma\left(\alpha + \underbrace{\alpha_2^m + \frac{1}{2}}_{\equiv\alpha_2},\, \beta + \underbrace{\beta_2^m + \frac{1}{2}(y_2 - \mu)^2}_{\equiv\beta_2}\right)$$

From $\beta_2^m|\,\omega \sim \Gamma(\alpha_2^m,\, \omega)$ and $\frac{1}{2}(y_2 - \mu)^2\Big|\,\omega \sim \Gamma\left(\frac{1}{2},\, \omega\right)$, we get $\beta_2|\,\omega \sim \Gamma(\alpha_2,\, \omega)$. The memory variable is defined as

$$m_3 = \beta_2\, u_3$$

where $u_3 \sim \mathcal{B}(\lambda\alpha_2,\,(1-\lambda)\,\alpha_2)$. Then again, we see that

$$\omega|\, m_3 \sim \Gamma\left(\alpha + \underbrace{\lambda\,\alpha_2}_{\equiv\alpha_3^m},\; \beta + \underbrace{\beta_2\,u_3}_{\equiv\beta_3^m}\right)$$

Recursively, we can then express any time t prior as

$$\omega|\, m_t \sim \Gamma\left(\alpha + \alpha_t^m,\; \beta + \beta_t^m\right)$$

where $\beta_t^m|\,\omega \sim \Gamma\left(\alpha_t^m,\,\omega\right)$ is held, and we can also express the posterior as

$$\omega|\, m_t, y_t \sim \Gamma\left(\alpha + \alpha_t,\; \beta + \beta_t\right)$$

where

$$\begin{aligned}\alpha_t &= \alpha_t^m + \frac{1}{2}\\ \beta_t &= \beta_t^m + \frac{1}{2}\left(y_t - \mu\right)^2\end{aligned}$$

Then, using the distribution of $\beta_t^m|\,\omega$, we see that $\beta_t|\,\omega \sim \Gamma\left(\alpha_t,\,\omega\right)$. Using the definition of the memory variable m_{t+1} in (6), we can also derive that $m_{t+1}|\,\omega \sim \Gamma\left(\lambda\,\alpha_t,\,\omega\right)$. Finally, we derive $\omega|\, m_{t+1}$ below.

$$\begin{aligned}\omega|\, m_{t+1} &\propto f\left(m_{t+1}|\,\mu,\omega\right)\; f\left(\omega|\,\mu\right)\\ &\propto \left(\omega^{\lambda\,\alpha_t}\,\exp^{-\omega\,m_{t+1}}\right)\left(\omega^{\alpha-1}\,\exp^{-\beta\,\omega}\right) \propto \Gamma\left(\alpha + \alpha_{t+1}^m,\; \beta + \beta_{t+1}^m\right)\end{aligned}$$

where

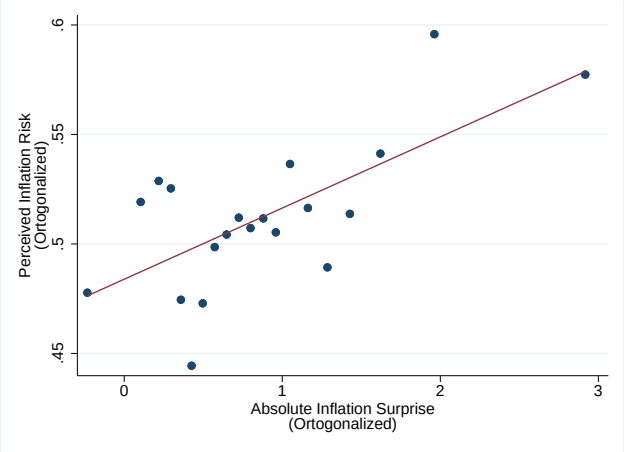
$$\begin{aligned}\alpha_{t+1}^m &= \lambda\,\alpha_t\\ \beta_{t+1}^m &= m_{t+1} = \beta_t\,u_{t+1}\end{aligned}$$

B Robustness Checks

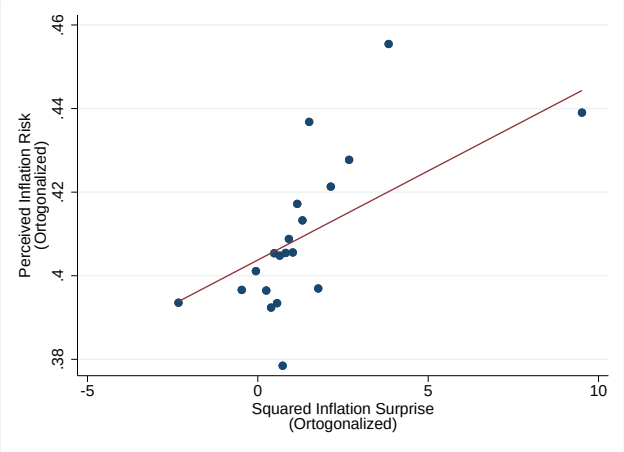
In figure 5, we show alternative versions of figure 4 in which we replace squared surprises with absolute revisions (panel (a)) or use Huber weights (panel (b)). The numerical estimates associated with these figures is presented in table 3

Figure 5: Scatter Plot of Baseline Regression: Robustness

(a) Absolute Surprises



(b) Huber-Robust



NOTE. Panel (a) shows a binned scatter plot of equation(14) (i.e., all variables have been orthogonalized to the controls in that regression) except that we have replaced the squared surprise with the absolute value of the surprise. Panel (b) shows estimates of our baseline specification using Huber weights, following Coibion et al. (2018).

Table 3: Empirical Estimation: Robusntess

x_{it}^2	0.0514	0.0352	
	(0.0218)	(0.0129)	
$\lvert x_{i,t} \rvert$			0.0346
			(0.0136)
Observations	6349	6136	6349
R^2	0.531	0.740	0.531
Sample	1992–2019	1992–2019	1992–2019
Specification	Baseline	Huber	Abs. Rev.

NOTE. This table shows estimates of equation (14). The first column is a replication of the first column (our baseline) of table 2—see the note to that table for more details. In the second column, we use Huber weights to eliminate the role of outliers. In the third column, we replace squared surprises with absolute surprises.