IIMB-WP N0. 729/2026

WORKING PAPER NO: 729

Long-Run Inflation Expectations, Bounded Rationality, and Inflation Uncertainty

Arnab Biswas
Assistant Professor,
Department of Finance,
IFMR GSB, Krea University,
arnab.biswas@krea.edu.in

Chetan Subramanian
Professor,
Department of Economics,
Indian Institute of Management Bangalore
chetan.s@iimb.ac.in

Year of Publication – January 2026

Long-Run Inflation Expectations, Bounded Rationality, and Inflation Uncertainty

Arnab Biswas[1] **Chetan Subramanian**[2]

This Version: December 20, 2025

Abstract

This paper studies the formation and anchoring of long-run inflation expectations under bounded rationality and adaptive learning. The empirical analysis applies this framework to India, which provides a useful setting given its recent transition to a formal inflation-targeting regime. Motivated by the persistence of elevated term premia in India's sovereign bond market, we develop and estimate a New Keynesian model in which agents update beliefs about long-run inflation in response to short-horizon forecast errors, implying that expectations may remain imperfectly anchored even under an explicit inflation-targeting regime. The model is estimated using inflation data and survey-based short-run inflation expectations from India, allowing us to recover quarterly movements in ten-year-ahead inflation expectations.

We find that the introduction of inflation targeting lowered the level of long-run inflation expectations and moved them closer to the policy target. However, long-run beliefs remain sensitive to short-run inflation surprises, consistent with adaptive learning and incomplete anchoring. This persistent belief updating generates long-horizon inflation uncertainty, which is priced into nominal bond yields and contributes to elevated term premia at longer maturities. Counterfactual simulations show that fully anchored expectations would substantially reduce long-horizon inflation uncertainty and lower long-maturity term premia by around 120 basis points. The results highlight how bounded rationality in expectation formation can sustain inflation risk and term premia despite improvements in monetary policy credibility.

Keywords: Inflation Expectations, Bounded Rationality, Adaptive Learning, Bond Term Premium

JEL Classification: D83, D84, E42, E43, E44

[1]Arnab Biswas, Assistant Professor, IFMR GSB, Krea University, Department of Finance. Email: arnab.biswas@krea.edu.in

[2]**Corresponding author.** Chetan Subramanian, Professor, Indian Institute of Management Bengaluru, Department of Economics. Email: chetan.s@iimb.ac.in

1. Introduction

Monetary policy frameworks that rely on explicit inflation targets are designed not only to stabilize realized inflation but also to anchor long-run inflation expectations. When expectations are firmly anchored, temporary shocks to inflation have limited effects on long-horizon beliefs, allowing central banks to stabilize real activity without jeopardizing price stability (Bernanke, 2007; Woodford, 2003). However, if agents form expectations under bounded rationality and update beliefs through adaptive learning, long-run inflation expectations may remain sensitive to short-run inflation outcomes even in the presence of a credible numerical target.

This paper examines the anchoring of long-run inflation expectations within the framework of bounded rationality and adaptive learning. We develop and estimate a New Keynesian model in which agents update their perceived long-run inflation target in response to short-horizon inflation forecast errors. The empirical analysis applies this framework to India, which provides a useful setting given its recent transition to a formal inflation-targeting regime. In February 2015, the Government of India and the Reserve Bank of India adopted a 4% CPI inflation target with a tolerance band of 2–6%, creating a clear numerical anchor against which the evolution of long-run expectations can be evaluated.

Despite this institutional shift, financial-market outcomes suggest that long-run inflation uncertainty may not have been fully resolved. Figure 1 and the Table 1 documents a divergence in the Indian government bond market following the adoption of inflation targeting. The expectations-hypothesis component of the zero-coupon yield curve declines sharply across maturities, consistent with lower expected future short-term rates. In contrast, the term-premium component remains elevated and persistent, particularly at longer maturities. This coexistence of declining expected rates and sustained term premia points to residual long-horizon inflation risk, providing a natural motivation to study whether long-run inflation expectations remain imperfectly anchored and continue to respond to short-run inflation surprises.

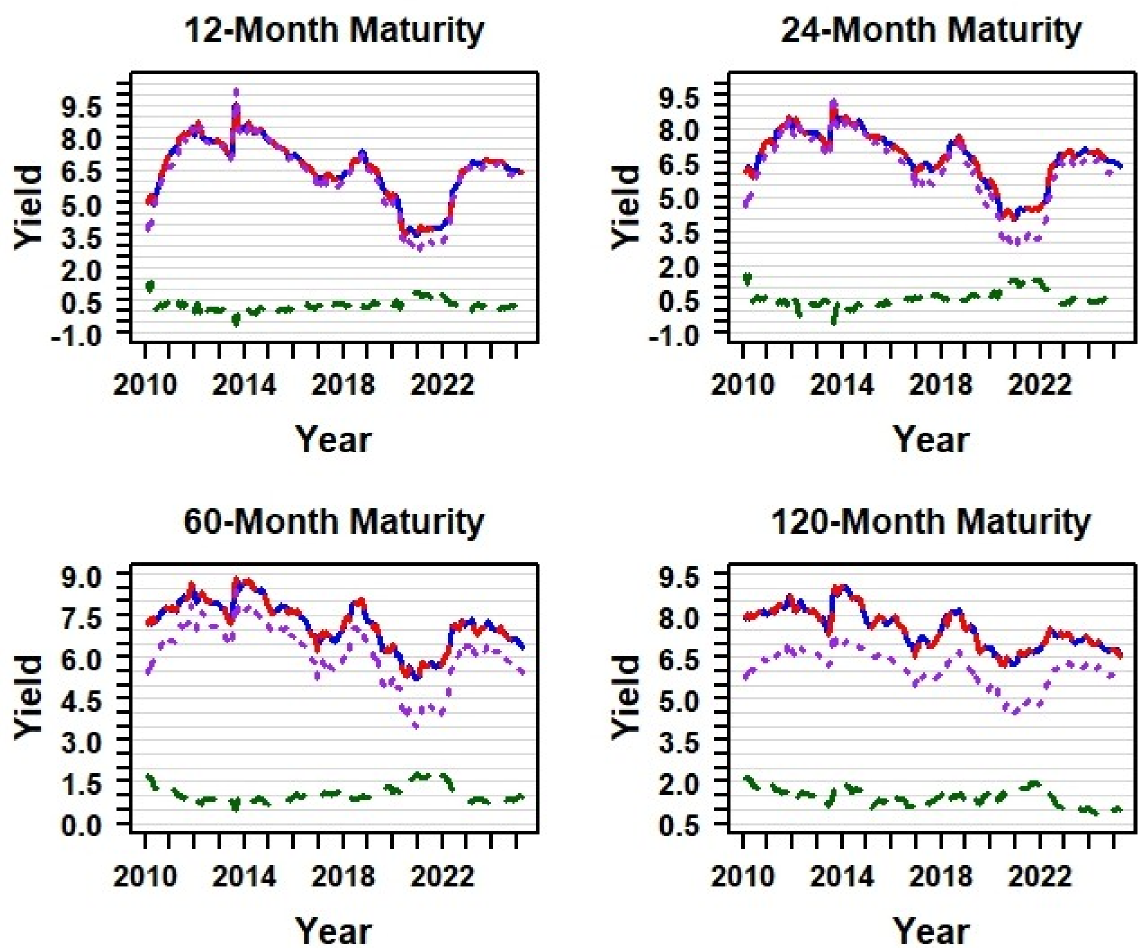

Figure 1: **ZCYC Yield and Term Premium at a given maturity: 12-, 24-, 60-, and 120-month maturities :**Each panel reports monthly yields and corresponding term premia over the period Q1:2010 to Q4:2024.The red and blue lines denote the observed and model-fitted yields, respectively. The purple dotted line represents the expectations-hypothesis component, while the green dotted line displays the estimated term-premium component.

Table 1: **Term Premium Estimates Across Maturities**

Maturity (months)	**Window**	**Average Min** (%)	**Average Max** (%)	**Sample Mean** (%)
12	Pre-IT	-0.198	0.540	0.179
12	Post-IT	0.156	0.524	0.345
24	Pre-IT	-0.050	0.667	0.328
24	Post-IT	0.456	0.868	0.663
60	Pre-IT	0.734	1.128	0.934
60	Post-IT	0.985	1.308	1.128
120	Pre-IT	1.338	1.848	1.599
120	Post-IT	1.175	1.848	1.599

Notes: The Table depicts Term premium estimates across ZCYC maturities. Pre-IT corresponds to the period Q1:2010 to Q1:2015, and post-IT is Q2:2015 to Q4:2024. Term premium estimates are on a monthly basis and estimated following Adrian et al. (2015) and corrected for small-sample bias via inverse bootstrapping. The sample is divided into 12-month blocks. For each block, the minimum and maximum term premia are computed; the reported "average minima" and "average maxima" are the means of these block-wise values.

This divergence suggests that improvements in short-run inflation performance need not translate into fully anchored long-run expectations, particularly when agents update beliefs based on realized inflation outcomes. The analysis is organized around three issues. First, how have long-run inflation expectations evolved following the adoption of inflation targeting, as measured by ten-year-ahead expectations? Second, to what extent are long-horizon expectations anchored to the announced 4% target, and how sensitive are they to short-run inflation surprises? Third, does residual long-run inflation uncertainty contribute to the persistence of term premia at longer maturities of government securities?

Empirical analysis of long-term inflation-expectation anchoring in India is constrained by the absence of market-based indicators. Unlike advanced economies that rely on instruments such as TIPS breakevens or inflation swaps across maturities, India offers no comparable measures from which long-horizon inflation expectations can be directly inferred. To address this limitation, we conduct our analysis within a New Keynesian framework following Carvalho et al. (2023), featuring a central bank that credibly commits to a fixed inflation target under a time-invariant policy regime. Firms operate under monopolistic competition and nominal rigidities, setting prices optimally based on expectations of future marginal costs and the perceived inflation target, with marginal costs determined by monetary policy and aggregate demand conditions.

A central feature of the model is the adaptive learning process through which agents form beliefs about the long-run inflation target, generating endogenous uncertainty around the expected target. Within this learning environment, agents exhibit bounded rationality: long-run expectations need not be fully anchored and are revised in response to short-run forecast errors in realized inflation. As a result, learning dynamics create a wedge between agents' long-term inflation expectations and the fixed policy target. Expectation formation follows the switching-gain algorithm of Carvalho et al. (2023), which governs the degree of anchoring of long-run beliefs. Under a constant-gain specification, expectations remain backward-looking and unanchored, whereas under a decreasing-gain specification, expectations gradually converge toward the policy target in a forward-looking manner. Consistent with Bernanke (2007), anchoring thus emerges endogenously from the interaction between short-run inflation surprises and agents' long-run belief updating

We estimate the model using headline CPI inflation data from 1958 to 2024 and compare the model-implied long-run inflation expectations with survey-based inflation expectations from the Reserve Bank of India's Survey of Professional Forecasters (SPF). Specifically, we use one- to four-quarter-ahead SPF expectations for the period 2010–2024 to examine the evolution of expectations before and after the adoption of inflation targeting. Direct measures of long-horizon expectations are limited in the Indian context: the only available source of ten-year-ahead inflation expectations is the RBI's SPF, which is available for a short window from Q1:2008 to Q1:2018. Over this overlapping period, the model-implied long-run expectations closely track the survey-based ten-year-ahead forecasts, while allowing us to extend estimates of long-horizon expectations through the end of the sample period.

The results of this paper indicate that long-term inflation expectations in India declined notably following the adoption of the 4% inflation target. In the pre-IT era, when policy communication was less explicit and information asymmetries were substantial, long-horizon expectations were highly responsive to short-term inflation surprises. Post-IT, expectations have moved closer to the stated target: by Q4:2024, the model places ten-year inflation expectations within the 4.0–4.7% range (99% posterior interval). However, the underlying learning behavior suggests that long-run expectations are not yet firmly anchored. Although the median sensitivity to transitory inflation shocks has declined, the distribution of this sensitivity remains skewed to the upside. The intuition is straightforward: in an environment where inflation repeatedly approaches or breaches inferred "rational" bounds, in our case the "upper" rational bounds are frequently breached, agents continue to assign disproportionate weight to the risk of upward deviations. This is consistent with the model-implied rational range of approximately 3.8–4.9%, which is considerably narrower than the central bank's $\pm 2\%$ operational tolerance band. The gap between the rational bounds and the official range underscores the need for sustained inflation stabilization if long-run expectations are to converge decisively toward the announced target

A central insight from the adaptive-learning framework is that credibility accumulates only when inflation remains within the rational bounds for an extended period—on the order of several years. India has not yet achieved such stability. Even in the post-IT period, inflation has frequently hovered near the upper rational bound and has occasionally exceeded it, often during cyclical slowdowns. These episodes, though smaller than the large deviations seen before IT, occur regularly enough to disrupt the transition toward a decreasing-gain, forward-looking learning regime. Consequently, agents continue to update long-run expectations based on near-term inflation outcomes, reflecting a form of persistent backward-looking behavior despite the presence of an explicit numerical target. Finally, the decline in long-run expectations after IT reflected not only the policy shift but also the behavior of the inflation process itself. The early IT period coincided with subdued non-core inflation and a broader cyclical slowdown, generating a series of negative inflation surprises. Because agents remained partially backward-looking when IT was introduced—carrying forward sensitivities from the pre-IT regime—these negative surprises had a disproportionate effect on long-run beliefs. A counterfactual imposing an immediate shift to a decreasing-gain regime shows that, had expectations been fully anchored at the outset, the decline in long-run inflation expectations between 2016 and 2020 would have been considerably smaller. Fully anchored expectations would not have reacted so strongly to transitory negative surprises.

These dynamics carry implications for bond pricing. Countercyclical breaches of the rational bounds raise inflation uncertainty precisely when economic conditions are weak, generating long-run risk that is particularly costly for investors. Given their preference for early resolution of uncertainty, investors demand compensation for bearing this risk, which manifests as a persistent term premium at longer maturities. Consistent with this mechanism, the expectations-hypothesis component of the yield curve declined sharply after 2015—bringing ten-year yields from roughly 9% in 2014 to around 6.7% by 2021—yet the term-premium component remained largely unchanged, typically in the 100–200 basis-point range. The upside skew in forecasters' sensitivity to transitory inflation disturbances provides a direct mechanism for this persistent premium.

A complementary long-run-risk asset-pricing model following Song (2017) and (Bansal and Shaliastovich, 2013; Bansal and Yaron, 2004) reinforces these findings. Allowing perceived inflation-target beliefs to drift in response to shocks—particularly those that are countercyclical—enables the model to match the observed persistence and magnitude of term premia. In a counterfactual environment characterized by strict inflation stabilization around the policy target, analogous to an environment of complete anchored expectations, the covariance channel that drives nominal long-run risk is muted, resulting in a roughly 120-basis-point decline in the term premium. This contrast strengthens the central conclusion that long-term expectations in India are not yet fully anchored: the possibility of drift in the perceived steady state remains a priced source of risk.

Contribution to the Literature:

In the Indian context, several studies have examined the effects of inflation targeting on inflation expectations. Eichengreen et al. (2021) report declines in realized inflation and reductions in both household and Survey of Professional Forecasters (SPF) short-term expectations (one to four quarters ahead) in the post-IT period, accompanied by lower persistence and volatility. These patterns are interpreted as evidence of improved short-run anchoring. Using household survey data, Asnani et al. (2019) similarly document improved anchoring following the adoption of IT. Their analysis focuses on the distribution of qualitative inflation expectations across five response categories, ranging from expected price declines to increases above the current rate. It relates these to realized inflation outcomes and core and non-core price volatility. They find that roughly 31 percent of respondents move from expecting high inflation to more moderate or stable outcomes after IT, which the authors interpret as enhanced anchoring, despite limited changes in non-core inflation volatility.

More recent work reinforces these findings. Pattanaik et al. (2023) construct an inflation-expectations anchoring index using aggregate household survey data and show stronger anchoring performance in the post-IT regime. Garga et al. (2022) examine financial-market expectations based on OIS curve , and conclude that the adoption of IT was viewed as a credible policy shift, as reflected in a more systematic monetary policy response to inflation. Kishor and Pratap (2023) decompose one- to four-quarter-ahead headline inflation expectations into trend and cyclical components following Stock and Watson (2007). They find that, after IT, the trend component becomes less sensitive to negative sentiment shocks and the cyclical component exhibits lower volatility, both of which are interpreted as improved short-term anchoring.

A significant limitation of the existing literature on India is its focus on short-term expectations. Most studies examine forecasts for one to four quarters ahead. However, the widely accepted benchmark for anchoring expectations, established by Bernanke (2007), emphasizes the long-term behavior of expectations and their response to temporary shocks. Consequently, we have a limited understanding of whether India's inflation-targeting framework has effectively stabilized expectations over longer horizons, where credibility is crucial. Interpreting the apparent improvement in short-term expectations is also complicated by the macroeconomic environment of the early post-IT years. The transition to IT coincided with a period of favourable price developments, particularly in food and oil, which together account for roughly 60% of India's CPI basket (Figures A1-A2). As shown by Ascari et al. (2016) and Ball and Mankiw (1995), such favourable shocks tend to pull down both the mean and the skewness of the inflation distribution. When this happens, short-horizon expectations naturally look less volatile, more persistent, and more symmetric—not because agents have become more firmly anchored, but because the shocks hitting the economy are unusually benign. Once these shocks reverse, the apparent stability typically dissipates.

Our gain estimates help separate this compositional effect from genuine anchoring. While the median sensitivity to inflation surprises has declined since the introduction of IT, the upper tail of the distribution remains elevated. This suggests that a significant proportion of agents continue to update their long-run expectations whenever inflation deviates from recent trends. In other words, some progress has been made, but a substantial portion of the forecasters remains cautious, still relying on incoming data rather than fully trusting the announced target.

Taken together, the evidence points to partial rather than complete anchoring. Short-term expectations may appear more stable, but much of that stability reflects favourable shocks rather than deep changes in belief formation. Long-run expectations—the true test of anchoring—continue to display sensitivity to temporary inflation movements, suggesting that credibility, while improving, has not yet fully taken hold. Our work contributes to a broad literature examining how information frictions and heterogeneous perceptions of the long-run inflation anchor shape expectation dynamics. Prior studies—such as (LeBaron and Smith, 2025; Malmendier and Nagel, 2016; Branch and Evans, 2011; Eusepi and Preston, 2010; Lansing, 2009; Preston, 2008)—show that when agents hold differing views about the steady-state inflation level, forecast errors can become self-reinforcing, generating persistent drift in long-horizon expectations. In such environments, consistent and credible policy conduct is critical for stabilizing beliefs.

Our paper also relates to the literature on adaptive learning and gradual belief updating (Coibion et al., 2018; Guse, 2014; Milani, 2007; Honkapohja and Mitra, 2003; Evans and Honkapohja, 2003). In these frameworks, agents revise their perceived long-run inflation target based on observed outcomes, generating inertia and persistence in expectations that align with key empirical features of macroeconomic time series. Similar learning dynamics can produce deviations of asset prices from fundamentals, giving rise to endogenous risk premia, as shown by Kim (2009) and Chakraborty and Evans (2008). However, this body of work focuses primarily on advanced economies. Our contribution is to examine how long-run inflation targets drift in both pre- and post-inflation-targeting periods in an emerging market context.

2. Long-term inflation expectations

This section first introduces the benchmark New Keynesian framework, then details the adaptive-learning mechanism integrated into its benchmark framework following Carvalho et al. (2023). The benchmark economy features monopolistic competition, with firms maximizing profits and the central bank focused on stabilizing inflation around a set target. Information asymmetry is captured by a mechanism that endogenously determines rational bounds on agents' short-term inflation expectations and innovations in their evolving beliefs about the long-run inflation target and expectations. Belief updates are triggered by surprises in realized inflation relative to each agent's rational bounds, linking short-term forecast errors to adjustments in long-horizon beliefs. These evolution in long-run target beliefs, in turn, prompt revisions in short-term forecasts, creating a feedback loop. The model is estimated using Bayesian methods, generating distribution of time-varying long-term inflation expectations, short-term rational bounds, and the distribution of associated sensitivity (gain function) measures to short-term forecast errors.

2.1. The Benchmark Model: Firm Price-Setting Problem

A continuum of monopolistically competitive firms maximizes discounted profits, subject to Rotemberg (1982)-type adjustment costs. The optimization problem for each firm i is as follows:

$$E_t \sum_{T=t}^{\infty} (\gamma\beta)^{T-t} M_{t,T} \left[Y_T(i) \left(\frac{P_t(i)}{P_T} - mc_t \right) \right] \tag{2.1}$$

In this context, $M_{t,T}$ denotes the stochastic discount factor, γ represents the stable eigenvalue of the model, which is analogous to the Calvo probability that prices remain unchanged each period, and β is the intertemporal discount factor. The demand curve for firm i is given by $Y_t(i) = \left(\frac{P_t(i)}{P_t}\right)^{-\theta_{p,t}} Y_t$, where $\theta_{p,t}$ is the elasticity of substitution across varieties. In this formulation, $P_t(i)$ is the price set by firm i relative to the aggregate price level P_t, and mc_t denotes real marginal cost. The optimal pricing rule, conditional on a fixed inflation target π^*, is expressed as follows:

$$P_t(i) = E_t \sum_{T=t}^{\infty} (\gamma\beta)^{T-t} \big[(1-\gamma\beta)(mc_t + u_T) \big] + \gamma\beta \big((\pi_{T+1} - \pi^*) - \gamma_p (\pi_T - \pi^*) \big) \tag{2.2}$$

where γ_p is the degree of indexation to lagged inflation, and u_T represents cost-push shock. Log-linearizing the pricing condition and assuming symmetry across firms yields the aggregate supply relation in log-deviation form:

$$\hat{\pi}_t = \gamma_p \hat{\pi}_{t-1} + \mu_t + E_t \sum_{T=t}^{\infty} (\gamma\beta)^{T-t} \big[\tau_p \widehat{mc}_t + (1-\gamma)\beta(\hat{\pi}_{T+1} - \gamma_p \hat{\pi}_T)\big] \tag{2.3}$$

Here, $\tau_p = (1-\gamma)(1-\gamma\beta)/\gamma$, $\hat{\pi}_t$ represents deviations of inflation from the steady state, and $\mu_t \sim N(0,1)$ denotes the exogenous cost-push disturbance.

2.1.1. The Benchmark Model: Discretionary Monetary Policy

The benchmark model incorporates a Central Bank that stabilizes inflation around a fixed target π^* and implements a discretionary stabilization policy, similar to Woodford and Walsh (2005), presented in log-linearized form:

$$\hat{\pi}_t - \gamma_p \hat{\pi}_{t-1} + \Gamma_x x_t = \epsilon_t. \tag{2.4}$$

x_t represents the output-gap deviation. The exogenous shock ϵ_t, resulting from discretionary stabilization policy, is persistent and follows an autoregressive process of order one (AR(1)):

$$\epsilon_t = \rho\, \epsilon_{t-1} + \xi_t, \qquad \xi_t \sim \mathcal{N}(0,1).$$

The exogenous shock can be interpreted as an error arising mainly from persistent policy lapses and mismeasurement of inflation and output deviations from their respective steady states. Marginal-cost deviations are assumed to be proportional to the output gap, $\widehat{mc}_t = \phi_t x_t$, and for analytical tractability and to facilitate the selection of priors, ϕ_t is set to unity.

2.1.2. The Benchmark Model: Rational Expectations of the Representative Agent

Equation (2.3) may be reformulated into the standard forward-looking New Keynesian specification:

$$\hat{\pi}_t - \gamma_p \hat{\pi}_{t-1} = \tau_p \widehat{mc}_t + \beta E_t(\hat{\pi}_{t+1} - \gamma_p \hat{\pi}_t) + \mu_t. \tag{2.5}$$

By substituting equations (2.4) and (2.5), and applying $\widehat{mc}_t = x_t$, the representative agents' stationary rational-expectations representation for one-period-ahead inflation in log deviation form is obtained:

$$\hat{\pi}_t = \gamma_p \hat{\pi}_{t-1} + \rho\,\bar{\omega}\,\epsilon_{t-1} + \nu_t, \tag{2.6}$$

where $\bar{\omega} = \left[\,1 + (1-\beta\rho)\tau_p^{-1}\Gamma_x\,\right]^{-1}$ and $\nu_t = \bar{\omega}\,\xi_t + \left(1 + \tau_p \Gamma_x^{-1}\right)^{-1}\mu_t$. This formulation highlights the role of policy inertia, represented by ρ and the composite weight $\bar{\omega}$, in transmitting shocks ξ_t and μ_t to near-term inflation dynamics under indexation γ_p.

2.1.3. The Perceived Law of Motion

This subsection and those that follow detail the integration of the adaptive framework into the benchmark model. The baseline benchmark model assumes that π^* is non-stochastic. However, agents may doubt the central bank's ability to strictly maintain inflation within rational bounds around the announced target, π^*. As a result, agents incorporate a time-varying perceived target $\bar{\pi}_t$ into their forecasting models. Under the assumption of symmetric expectations regarding $\bar{\pi}_t$, the stationary rational-expectation expression in (2.6) is modified to include $\bar{\pi}_t$ for near-term inflation forecasts. This modification yields a representative Perceived Law of Motion (PLM), expressed as:

$$\textbf{PLM} : \hat{\pi}_t = \gamma_p \hat{\pi}_{t-1} + (1-\gamma_p)\bar{\pi}_t + \rho\,\bar{\omega}\,\epsilon_{t-1} + e_t. \tag{2.7}$$

In this context, e_t denotes the one-step-ahead forecast error, which corresponds to the inflation surprise. Within the PLM, agents account for both the immediate past inflation deviation $\hat{\pi}_{t-1}$ from the fixed target and the perceived deviation in the inflation target$\bar{\pi}_t$ itself. Consequently, any drift in the target deviation $\bar{\pi}_t$ directly revises one-period-ahead inflation expectations and, through recursion, shapes expectations at longer horizons. At this stage, two key assumptions are introduced to connect short-run forecasts and surprises with adaption in long-run beliefs. The first assumption asserts that, in each period, long-run inflation expectations are equivalent to the perceived steady-state inflation target:

$$\lim_{T\to\infty} \hat{E}_t \pi_T \;=\; \bar{\pi}_t.$$

The second assumption posits that, under perfect foresight in the short run, long-run beliefs are updated each period based on realized inflation and short-run surprises observed up to the immediate past:

$$\hat{E}_{t-1}\bar{\pi}_T \;=\; \bar{\pi}_t.$$

Taken together, these assumptions indicate that changes in long-run inflation beliefs in each period, represented by $\bar{\pi}_t$, are conditional on short-term surprises. Through the PLM, these changes in long-run beliefs are incorporated into subjective near-term inflation forecasts.

2.2. The Actual Law of Motion

Within this framework, each firm projects the future trajectory of marginal costs based on its expectations of long-run inflation target and the degree of monetary-policy inertia. Under symmetric conditions, utilizing the discretionary monetary policy rule and by incorporating the time-varying deviations in the inflation target, the projected path of marginal cost is expressed as follows:

$$\hat{E}_t \widehat{mc}_{T+1} = \frac{1}{\Gamma_x} \hat{E}_t \left[\epsilon_{T+1} - (\hat{\pi}_{T+1} - \bar{\pi}_T) + \gamma_p (\hat{\pi}_T - \bar{\pi}_T) \right] \tag{2.8}$$

Here, monetary policy inertia is characterized by $\hat{E}_t \epsilon_{T+1} = \rho^{T-t} \epsilon_t$.

In each period, firms utilizes the extrapolated marginal cost path to re-optimize prices according to equation (2.1) subject to their adjustment cost. Aggregating across all the firms yields the representative Actual Law of Motion (ALM) for realized inflation in stationary form:

$$\textbf{ALM} : \hat{\pi}_t = \gamma_p \hat{\pi}_{t-1} + (1 - \gamma_p) \Gamma \bar{\pi}_t + \rho \bar{\omega} \epsilon_{t-1} + \nu_t \tag{2.9}$$

$$\text{where, } \Gamma = \frac{1}{1 + \tau_p \Gamma_x^{-1}} \frac{(1-\gamma)\beta}{1 - \gamma\beta}$$

The ALM represents the underlying inflation process, and is latent due to subjective expectations of each agent, as described by their respective PLM. Under symmetric aggregate assumptions, the gap between ALM- and agents' PLM-implied expectations, defined as the difference between equations (2.9) and (2.7), is primarily determined by the term $\Gamma \bar{\pi}_t$, where $\Gamma < 1$. The coefficient Γ quantifies the aggregate feedback from deviations in the inflation target beliefs to deviations in realized inflation. Thus, Γ provides a summary measure of the central bank's effectiveness in stabilizing inflation around its policy target in each period.

When the central bank adopts a policy path characterized by strict inflation targeting, the weight on the output gap coefficient (Γ_x) declines, which subsequently reduces Γ, as indicated in equation (2.9). This reduction in Γ diminishes the pass-through from deviations in the perceived inflation target $\bar{\pi}_t$ to realized inflation deviations $\hat{\pi}_t$. Strict inflation targeting also causes deviations in $\bar{\pi}_t$ to approach zero. In the limiting case where $\Gamma \to 0$, the ALM and PLM converge to the benchmark rational expectations around a non-stochastic steady-state inflation target, as described in equation (2.6). Consequently, a policy focused on strict inflation stabilization anchors long-term beliefs to the policy target, thereby eliminating drifts in perceived inflation target and their feedback to realized inflation.

Conversely, when monetary policy places greater emphasis on stabilizing the output gap rather than inflation, weightage of Γ_x increases in the policy path, resulting in a higher Γ. As a result, deviations in the perceived inflation target $\bar{\pi}_t$ exert a strong feedback effect on realized inflation deviations. In the extreme case where $\Gamma_x \to \infty$ (or equivalently $\Gamma \to 1$), the ALM converges with the PLM, and deviations in the perceived inflation target, are fully accounted for in realized inflation deviations. Therefore, as policy shifts away from active inflation stabilization, any changes in long-term beliefs about the inflation target become significant driver of innovations in realized inflation.

The empirical specification of the framework outlined in the subsequent subsections employs reduced-form systems that include simplified representations of both the PLM and the ALM. Within the system of equations, the coefficient term Γ_x does not appear explicitly. When calibrated to empirical data, the coefficient Γ in the ALM also reflects, in reduced form, deviations from strict inflation stabilization that arise from factors beyond output-gap stabilization.

The following subsection illustrates the update mechanism in perceived inflation target deviations, the self-fulfilling channel linking these revisions to realized inflation deviations, and the formulation of rational bounds.

2.3. Inflation Targeting Update Rule

The perceived inflation target deviations are updated each period as per the iterative mechanism:

$$\bar{\pi}_t = \bar{\pi}_{t-1} + g_t^{-1} e_{t-1} \tag{2.10}$$

where g_t, is the gain function and implies the weight assigned to the short-term inflation surprises e_t featuring in the PLM process (2.7). This gain function has it own period by period update rule, and exhibits switching between two functional form, which forms the basis of adaptive learning:

$$g_t = \begin{cases} g_{t-1} + 1, & \text{if } \left| \hat{E}_{t-1}\hat{\pi}_t - E_{t-1}\hat{\pi}_t \right| \leq \Theta\sqrt{\text{MSE}} \quad \textbf{(Decreasing gain)}, \\ \bar{g}^{-1}, & \text{otherwise} \quad \textbf{(Constant gain)}. \end{cases} \tag{2.11}$$

Here, $\hat{E}_{t-1}\hat{\pi}_t$ denotes the subjective forecast based on the PLM of the representative agent, whereas $E_{t-1}\hat{\pi}_t$ refers to the forecast implied by the ALM. The term $\sqrt{\text{MSE}}$ represents the standard deviation of the exogenous innovation ν_t in the ALM process (2.9), where $\nu_t \sim \mathcal{N}(0, \sigma_\nu)$. The threshold Θ defines the rational bound and characterizes the limits to rationality agents possess regarding the underlying inflation process. In this framework, switching in the gain function depends on the ratio of the difference between PLM and ALM based expectations normalized by the standard deviation of innovations in the ALM process, relative to the bound

Θ.

A lower ratio relative to Θ suggests that agents' PLM closely aligns with the ALM, indicating a high degree of forecast accuracy. When this ratio consistently remains within the bound Θ, the inflation process defined by the ALM is stable and predictable by the PLM. As a result, agents continue to employ decreasing gain learning and gradually reduce the weight assigned to inflation surprise e_t in their PLM. Under persistent decreasing gain learning, $\bar{\pi}_t \approx \bar{\pi}_{t-1}$, and the perceived inflation target converges to a fixed value over time, rendering updates in the perceived target negligible. Consequently, the feedback channel from inflation target deviations to realized inflation deviations is diminished. Agents thus exhibit more forward-looking behavior as they become structurally aware of the long-run steady state.

However, if the ratio persistently exceeds Θ, agents adopt constant-gain learning and assign a fixed weight to the surprise e_t while updating their beliefs about the long-run inflation target. In this regime, long-term beliefs evolve continually, as structural awareness is maintained through continual adjustments to $\bar{\pi}_t$. These beliefs become highly sensitive to short-term inflation surprises relative to the rational bound Θ, leading agents to exhibit more backward-looking behavior when updating their long-term beliefs.

The difference between the expectations implied by the ALM and PLM can be further represented in terms of accumulated past short-term forecast errors, as shown below:

$$
\begin{aligned}
\mid \hat{E}_{t-1}\hat{\pi}_t - E_{t-1}\hat{\pi}_t \mid &= \mid (1-\gamma_p)(\Gamma - 1)\bar{\pi}_t \mid \\
&= \mid (1-\gamma_p)(\Gamma - 1)(\bar{\pi}_0 + \sum_{T=0}^{t} g_T^{-1} e_T) \mid
\end{aligned} \tag{2.12}
$$

The expanded equation (2.12) shows that recurrent and substantial short-term surprises in previous periods can consistently keep the ratio in (2.11) above the bound Θ, thereby maintaining agents in the constant-gain regime. Because the expanded equation is in log-deviation form, the weights assigned to past inflation surprises decay exponentially, with greater emphasis on recent observations. As agents transition from a highly volatile to a less volatile regime, they gradually reduce the weight placed on immediate short-term inflation surprises, as they adapt to the new environment. Accordingly, the ratio in (2.11) ensures a gradual reduction in the weights assigned to recent inflation surprises. These results in gradual weakening of the innovations in the perceived inflation target deviations and its feedback to realized inflation.

In summary, sensitivity to inflation surprises within this framework decreases significantly only when inflation remains within rational bounds for an extended period. Under these conditions, empirical measures of switching sensitivity will also display low skewness, indicating low uncertainty about long-term beliefs among all agents in the economy.

2.4. Reduced- Form representation

The system of equations from (2.7) to (2.11) is presented below in reduced log-deviation form. This reduced set constitutes the latent-state-space system, which facilitates the estimation of evolution in the perceived inflation-target deviation from observable data.

$$\textbf{ALM} :: \hat{\pi}_t = (1-\gamma)\Gamma\bar{\pi}_t + \gamma_p\hat{\pi}_{t-1} + \epsilon_t + \mu_t$$

$$\textbf{Inflation-Target Update} :: \bar{\pi}_t = \bar{\pi}_{t-1} + g_t^{-1} e_{t-1}$$

$$\textbf{Forecast Error} :: e_t = (1-\gamma)(\Gamma - 1)\bar{\pi}_t + \mu_t + \varepsilon_t$$

$$\textbf{Monetary Policy Inertia} :: \epsilon_t = \rho_\epsilon \epsilon_{t-1} + \varepsilon_t$$

$$\textbf{Gain Update Rule} :: g_t = \begin{cases} g_{t-1} + 1, & \text{if } \left|\hat{E}_{t-1}\hat{\pi}_t - E_{t-1}\hat{\pi}_t\right| \leq \Theta\sqrt{\varepsilon_t^2 + \mu_t^2} \quad \textbf{(Decreasing gain)}, \\ \bar{g}^{-1}, & \text{otherwise} \quad \textbf{(Constant gain)}. \end{cases}$$

The variables $\hat{\pi}_t$, ϵ_t, ε_t, and μ_t constitute the linear state variables, whereas $\bar{\pi}_t$ and g_t represent the nonlinear state variables in the system. The shocks μ_t and ε_t are modeled as exogenous and are assumed to follow a standard normal distribution. The state vector comprising of non-linear state variables Δ_{1t} = ($\bar{\pi}_t$, g_t) and linear state variables Δ_{2t}=($\hat{\pi}_t$, ϵ_t, ε_t, μ_t) is mapped to observed data in levels via the measurement equation, as follows:

$$Y_t = \begin{bmatrix} \pi_t \\ E_t^{SPF}\pi_{t+1} \\ E_t^{SPF}\pi_{t+2} \\ E_t^{SPF}\pi_{t+3} \\ E_t^{SPF}\pi_{t+4} \end{bmatrix} = \pi_t^* + C_{1t}^{'}\Delta_{1t} + C_{2t}^{'}\Delta_{2t} + D_t o_t \tag{2.13}$$

Measurement errors, denoted by o_t, are assumed to follow a standard normal distribution. The intercept π_t^* aligns the model with observed variables in levels by capturing the mean long-run headline inflation rate. The observed headline realized inflation in levels is denoted by π_t, and the terms $E_t^{\text{SPF}}\pi_{t+h}$ for $h = 1, \ldots, 4$ correspond to one- to four-quarter-ahead headline inflation expectations.In total, fourteen structural parameters, as listed in the Table 2, are associated with the measurement and latent state space system and are estimated over the full sample. The measurement equation is recursively iterated to obtain model predictions of the inflation expectations across horizons. Appendix A provides details on the system of state equations, the parameters, and their relationship to observed variables through the measurement matrices C_t and D_t, and the recursions.

3. Methodology, Data, and Priors used in estimation

3.1. Calibration Strategies

As outlined in the preceding section, inflation dynamics within the reduced latent form state-space system are determined by both linear and nonlinear state variables.Estimation of this mixed state-space system employs sequential Monte Carlo methods, wherein the nonlinear states are marginalized via particle filters, as outlined by Schön et al. (2005). Particle filters facilitate the modeling of nonlinear belief evolution and endogenous transitions between the gain regimes. Conditional on each drawn set of nonlinear particles $\bar{\pi}_t$ and g_t, the linear state vector ($\hat{\pi}_t$, ϵ_t, ε_t, μ_t) is updated using the linear-state space Kalman filter. This methodology produces distributions for both linear and nonlinear filtered states, conditioned on the observed data. Importance sampling is conducted using the sequential Monte Carlo techniques of Andrieu et al. (2010) and Doucet et al. (2001). Parameter updates are performed via a Random-Walk Metropolis–Hastings (RWMH) algorithm with Robbins–Monro step-size adaptation to achieve optimal acceptance rates, following Vihola (2012) and Cai (2010). Based on posterior draws of the structural parameters and filtered states, smoothed state estimates are generated using the forward–backward sampling algorithm of Kitagawa (1993) and Carter and Kohn (1994), providing full time paths for posterior distribution of long-run inflation expectations, rational bounds, and the posterior distribution of the switching gain function.

Due to the limited availability of short- and long-horizon survey expectations for India, posteriors estimated from the comprehensive U.S. data set in Carvalho et al. (2023) are used as priors for all structural parameters except the mean inflation intercept, structural shock variances, and observation-error variances. This approach maintains well-defined rational bounds while allowing Indian data to inform the dynamics of state evolution and the sensitivity to inflation surprises. To balance the influence of prior information with domestic data, the in-sample likelihood for India is tempered using a scaling parameter $\lambda = 0.5$. This ensures that the posterior reflects domestic evidence without diverging substantially from the rational bounds established on the exhaustive U.S. sample. The tempered posterior is formally expressed as follows:

$$P^{In}(\bar{\Delta}^{In} \mid Y_t^{In}, Y^{US}, \bar{\Delta}^{US}) \propto L(Y_t^{In} \mid \bar{\Delta}^{US}, \bar{\Delta}^{In})^{\lambda}\, L(Y_t^{US} \mid \bar{\Delta}^{US})\, p(\bar{\Delta}^{US})\, p(\bar{\Delta}^{In}).$$

A total of 4,000 particles is used to approximate the nonlinear state space. The Random Walk Metropolis–Hastings (RWMH) algorithm is executed for 600,000 iterations with two independent chains. The last 100,000 iterations of each chain are reserved for the smoother, which is applied every ten steps. Step sizes are tuned to maintain an acceptance probability near 23.4%.

3.2. Data Source

Realized inflation is measured using the headline CPI, with quarterly data sourced from the CEIC International Monetary Fund series, and spans from the Q1:1958 to the Q4:2024. Inflation expectations, including short-term (one to four quarters ahead) and long-term (five and ten-year ahead), are sourced from the SPF conducted by the Central Bank of India. The SPFs are available on a quarterly basis from the Q1:2008 and have been conducted bimonthly since fiscal year 2014. From 2014 onwards, within-quarter averages are computed to construct a balanced panel for estimating the reduced-form equations. The SPF is the sole source of long-term inflation expectations (five and ten years ahead) in India and is available only for a limited period, specifically from Q1:2008 to the Q1:2018. To ensure compatibility with the SPF series, realized inflation for each quarter is defined as the change in the CPI index over the preceding four quarters. The moments of the realized inflation and the inflation expectation across horizons for both the pre and post IT regime are shown below in the Table 2.

Table 2: **Moments of Realized Inflation and Inflation Expectations before and after the Inflation Targeting Regime across horizons**

	Pre-IT		**Post-IT**	
	Mean	Std	Mean	Std
Realized Inflation (Full sample)	7.62	5.49	4.99	1.37
Realized Inflation (Q4:2007–Q1:2015)	9.21	2.51	–	–
1Q Ahead Expectation	8.65	2.02	4.93	1.13
2Q Ahead Expectation	8.11	1.83	4.83	0.79
3Q Ahead Expectation	7.64	1.50	4.76	0.65
4Q Ahead Expectation	7.29	1.26	4.67	0.60
5 Years Ahead Expectation	6.77	0.70	4.82	0.25
10 Years Ahead Expectation	6.07	0.67	4.57	0.24

Notes: One- to four-quarter-ahead expectations spans the period Q1:2008–Q4:2024. Five- and ten-year-ahead expectations are available for the period Q1:2008–Q1:2018.

Given the unbalanced nature of the dataset, filters are initially trained on realized inflation data from the Q1:1958 to the Q4:2007. Likelihood evaluation begins in the Q1:2008 and continues through the end of the sample period.

3.3. Prior and Posterior Estimates

Table 3 presents the priors and estimated posteriors. Given the infrequent and often implicit central-bank communication regarding the inflation objective during the pre–IT era, the prior for the intercept term π^* in the measurement equation is assigned a mean of 6.35%, consistent with the UCSV estimate of trend inflation for that period (following Stock and Watson (2007)).

Table 3: **Priors and Posterior Parameter Distributions of the Learning Algorithm**

	Prior			**Posterior**		
Parameter	**Distribution**	**Mean**	**Std. Dev.**	**Mode**	**1st Perc.**	**99th Perc.**
$\pi^*_{\text{post-IT}}$	Normal	4.35	0.10	4.45	4.13	4.57
$\pi^*_{\text{pre-IT}}$	Normal	6.35	0.10	6.33	6.08	6.54
θ	Gamma	0.022	0.006	0.011	0.010	0.060
$\bar{g}$	Gamma	0.140	0.029	0.410	0.110	0.490
Γ	Beta	0.906	0.041	0.930	0.540	0.990
ρ_ε	Beta	0.879	0.028	0.770	0.710	0.860
γ_p	Beta	0.126	0.028	0.380	0.470	0.490
σ_ε	IGamma	0.500	1.000	0.820	0.620	0.990
σ_μ	IGamma	0.500	1.000	0.960	0.680	0.990
σ_1	IGamma	0.500	1.000	0.740	0.250	0.990
σ_2	IGamma	0.500	1.000	0.590	0.390	0.860
σ_3	IGamma	0.500	1.000	0.380	0.270	0.650
σ_4	IGamma	0.500	1.000	0.380	0.330	0.740
σ_5	IGamma	0.500	1.000	0.470	0.450	0.940

Notes: The table reports prior distributions and posterior summary statistics for the parameters of the learning algorithm. IGamma denotes the inverse-gamma distribution.

For the post-IT era, the prior mean for π^* is set at 4.35%. The distribution of the π^* term encompasses both the pre- and post-IT sample means, as shown in Table 2. The parameters $\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5$ represent the standard deviations of the measurement errors for realized inflation and for one-to-four quarters-ahead SPF inflation expectations, respectively.

Both independent Markov chains achieve an acceptance rate of approximately 23.4%, and the multivariate Gelman test potential scale reduction factor is close to 1.07, suggesting proper mixing and convergence of the chains. The posterior estimates correspond closely with observed inflation dynamics. During the post-inflation-targeting (IT) period, the 99% confidence interval for the steady-state intercept π^* ranges from 4.24% to 4.52%, which is consistent with an average post-IT inflation rate of approximately 5%. The posterior mode of $\bar{g}$ suggests that, on average, agents in India assign greater weight to short-term forecast errors compared to their U.S. counterparts. The estimated degree of indexation, γ_p, exceeds its prior mean, indicating increased reliance on realized inflation in expectation formation. Similarly, the relatively high estimate of Γ implies a stronger self-referential channel than is typically observed in U.S. data. The standard deviations of the measurement errors are all below one, demonstrating that the model effectively captures the low-frequency movements in both realized inflation and inflation expectations.

4. Discussion

4.1. Benchmark Results

The learning algorithm effectively models the evolution of low-frequency inflation expectations across horizons. Quarterly variations in short-term expectations are tracked to a narrow margin, as demonstrated in Figures 2.

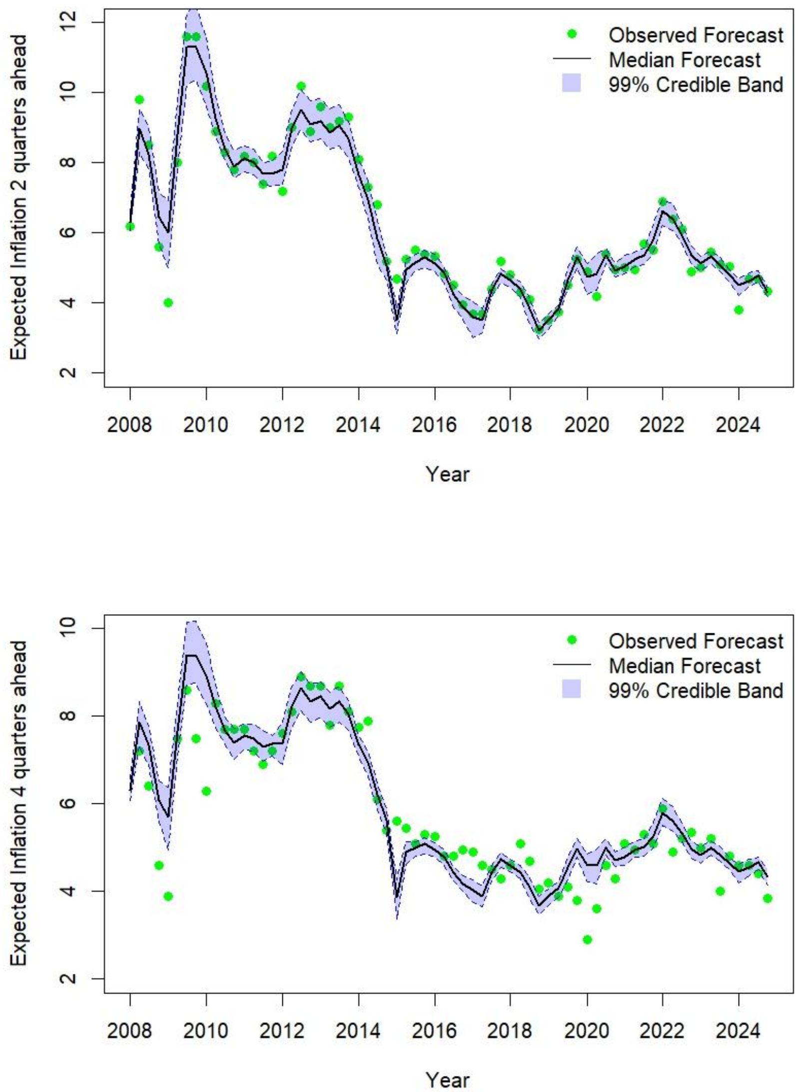

Figure 2: **Short-term inflation expectations for the period Q1:2007 to Q4:2024 :**The upper panel plots the model forecast together with observed two-quarters-ahead inflation expectations, while the lower panel shows the corresponding comparison for four-quarters-ahead expectations. Shaded areas denote the 99% posterior interval around the median forecast.

Feedback from deviations in the perceived long-term inflation target, combined with indexation to past inflation, is able to explain most of the variation in quarterly short-term inflation expectations. As indicated in Table 2 and indicated in the Indian literature, short-term inflation expectations declined significantly following the implementation of the inflation targeting (IT) regime and remained largely within the $4\% \pm 2\%$ operational bound. The volatility of short-term expectations also decreased substantially under the current IT regime. A significant driver of this reduction in the short-term inflation expectation volatility is the marked reduction in the drifts of the perceived long-term inflation target itself, along with the reduced volatility of the non-core CPI items (Figures in the Appendix A.1.3).

Evolution of the perceived long-term inflation target are assessed using the quarterly ten-year-ahead inflation expectations. Median model estimates closely track the survey-based ten-year-ahead expectations from the SPF, available through Q1:2018. A pronounced decline in expectations around 2015 aligns with the adoption of the 4% inflation target (Figure 3).

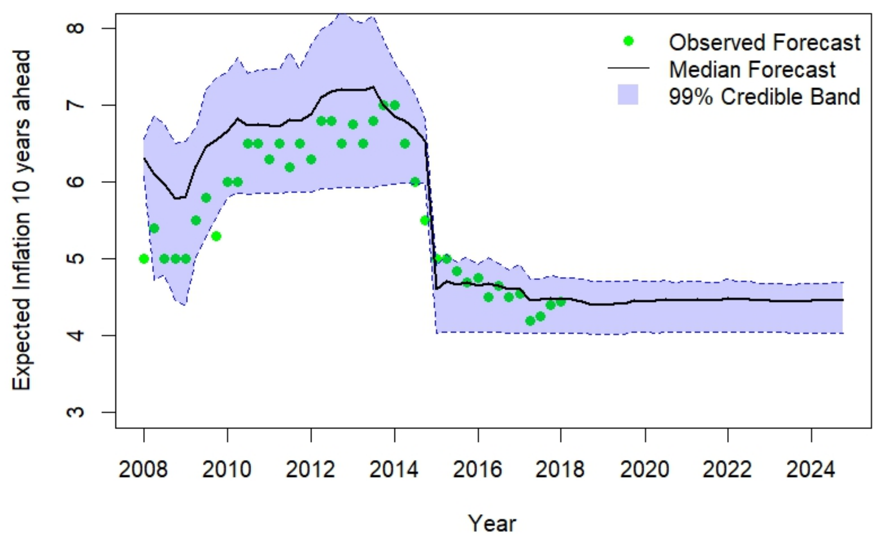

Figure 3: **Ten-year-ahead inflation expectations for the period Q1:2008 to Q4:2024 :**The figure plots the model forecast together with observed ten-year-ahead inflation expectations. Observed expectations are available for Q1:2008–Q1:2018; model-based forecasts extend to Q4:2024. Shaded bounds denote the 99% posterior interval around the median forecast.

Prior to the adoption of inflation targeting (IT), the lack of an explicit policy target resulted in significant drift in long-horizon expectations relative to the pre-IT average inflation rate of approximately 7%. Following the implementation of IT, this drift has substantially diminished. By the fourth quarter of 2024, ten-year-ahead expectations converge toward the 4% target, with 99% posterior estimates ranging from 4.0% to 4.7%. The confidence intervals are wider for ten-year-ahead expectations than for short-term expectations, reflecting the inherent higher uncertainty in long-term inflation forecasts.

The observed decline in ten-year-ahead expectations corresponds closely with the reduction in ten-year ZCYC yields, which is primarily attributable to the expectation hypothesis (risk-free) component in the post-IT regime (Figures 1 and 2). In the pre-IT regime, ten-year risk-free rates peaked at 7.39% in 2013–2014 and subsequently declined to a low of 4.48% during 2020–2021 in the post-IT regime. As the average real policy repo rate remains close to zero and is marginally negative in certain periods, the maxima and minima of risk-free yields closely follow average inflation in each regime: approximately 7.25% from the Q1:1958 to the Q4:2014 (pre-IT) and 4.52% from the Q1:2015 to the Q4:2024 (post-IT).

Although long-run expectations have improved with respect to to the policy target, uncertainty regarding the long-term inflation outlook persists. This ongoing uncertainty is demonstrated by the skewness of the sensitivity measure (switching gain) to short-term inflation surprises, as illustrated in Figure 4. While the median gain estimate declines following the adoption of IT, the posterior distribution of the sensitivity measure remains skewed.

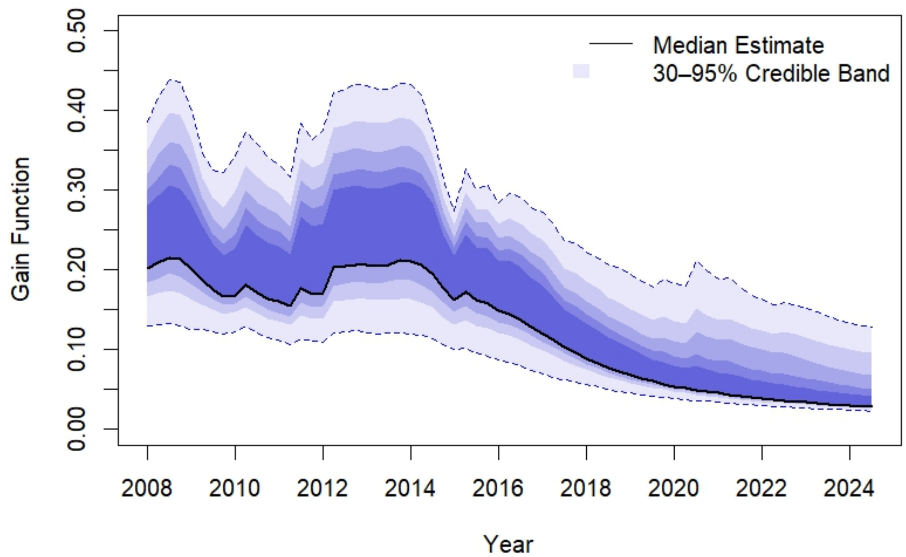

Figure 4: **Gain function for the period Q1:2008–Q4:2024 :**The figure displays the estimated gain function, with shaded areas denoting the posterior interval from the 30th to the 95th percentile.

This distribution represents the weight assigned by all agents in the economy to short-term inflation surprises. Within the representative-agent framework, where all agents are a priori exposed to the same information, persistent skewness indicates continued uncertainty about long-run inflation beliefs among a significant subset of agents. At higher posterior quantiles, the gain displays sharp switches at frequent intervals. Therefore, the upper tail of the distribution signals persistent, systematic inflation uncertainty, suggesting that a substantial fraction of agents remain backward-looking and that their expectations are not firmly anchored to the policy target. The skewness observed in the gain function arises mainly from frequent short-term violations of agents' estimated rational bounds for inflation outcomes, conditional on the perceived long-term inflation target. According to equations (2.11) and (2.12), the average value of the ratio

$$\frac{\Theta\,\sigma_\eta}{|(1-\Gamma)(1-\gamma_p)|} \approx 0.32,$$

and, given the post-IT 99% posterior range for π^*, the implied rational bound around the 4% policy target is approximately

$$3.8\% \ \leq \ \bar{\pi}_t + \pi^* \ \leq \ 4.9\%.$$

Maintaining inflation within this narrower bound promotes a symmetric, decreasing-gain learning regime and strengthens forward-looking expectations. Nevertheless, this bound is considerably narrower than the central bank's operational $\pm 2\%$ tolerance range, indicating that the central bank must implement more robust stabilization measures to ensure that agents' long-term inflation expectations remain anchored at the policy target.

Within the model, the weights attributed to short-horizon inflation surprises decay exponentially, becoming negligible beyond a five-year horizon. When inflation remains within the rational bound for an extended period (on the order of five consecutive years), the gain declines substantially, and the skewness in agents' sensitivity to inflation surprises is markedly reduced. In practice, however, inflation does not remain continuously within this bound. Even under the current IT regime, breaches of the upper rational bound persist in successive five-year periods. Although these breaches are less frequent and less pronounced than prior to the adoption of IT [1], they remain economically significant (Figure 5).

These episodes are predominantly counter-cyclical, frequently coinciding with OECD recession bars and downturns in the Composite Leading Indicator (CLI). Such periods correspond to spikes in the upper tail of the gain distribution, reflecting heightened long-run inflation uncertainty during economic downturns.

[1] The rational bound for the pre IT regime is estimated to be approximately in the range of 6% to 7%

This recurring pattern suggests that during periods of adverse growth, the effective anchoring of expectations systematically weakens, thereby amplifying the influence of short-run inflation surprises on long-run beliefs, even when an explicit IT regime is in place. As these episodes continue to occur in successive five-year periods, a significant fraction of agents persists in placing weight on these recurrent surprises while adapting their long-term beliefs.

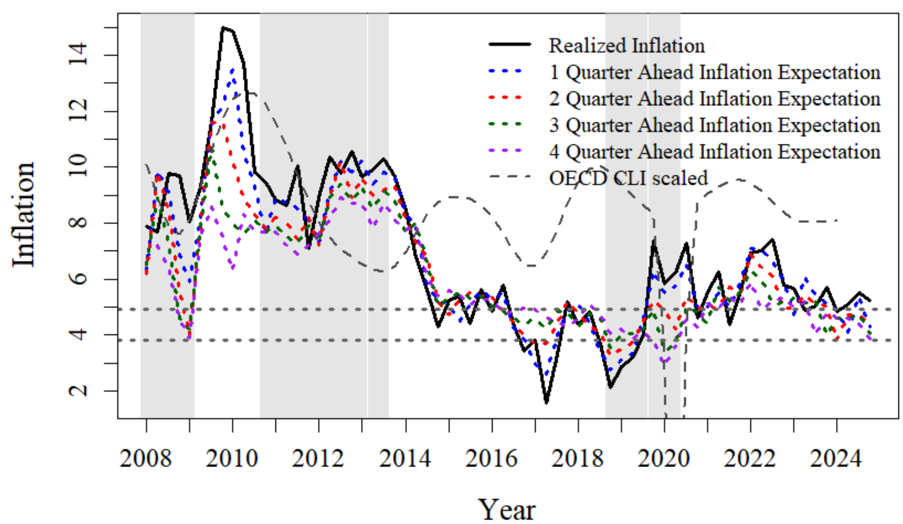

Figure 5: **Realized inflation and one- to four-quarter-ahead inflation expectations for the period Q1:2008–Q4:2024 :**Two horizontal dotted lines mark the rational bounds ranging between 3.8% to 4.9% in the IT regime. Shaded areas denote OECD recession bars (discontinued after 2022), and the solid black dotted line plots the OECD Composite Leading Indicator (CLI).

4.2. Counterfactual Estimates

This subsection first presents a counterfactual analysis to quantify how stabilizing inflation within the rational bound reduces the skewness of the switching gain function. A counterfactual inflation series is generated for the post-IT period to approximate an environment with more firmly anchored expectations and substantially reduced short-run deviations from rational bounds.

The counterfactual series is constructed using a multivariate Gaussian framework calibrated to represent the joint dynamics of realized inflation and one- to four-quarter-ahead expectations observed in the post–Q1:2015 sample. The empirical mean vector μ is estimated from this period, capturing the average values of realized inflation and short-term expectations across the specified horizons. To promote stability in the counterfactual regime, the associated covariance matrix Σ is rescaled by a shrinkage factor of 0.4, reducing covariance values to 16% of their empirical levels while preserving the original correlation structure. Concurrently, the mean of realized inflation is set at 4.45%, corresponding to the midpoint of the empirically estimated near-rational interval of 3.8% to 5.1%. The means of the one- to four-quarter-ahead expectations are adjusted downward from 4.45% by applying the historical differentials between realized inflation and each respective short-term expectation. This procedure maintains the historical dispersion between realized and expected inflation while shifting the joint distribution to a lower, more anchored regime.

A linear Gaussian transformation is applied to each post–Q1:2015 observation to map the observed data into the counterfactual regime: $y_t = \mu^{counter} + \nabla \left(x_t - \mu\right)$. In this formulation, $(\mu^{counter})$ denotes the target mean vector with realized inflation and expectations adjusted as previously described, and the matrix (∇) is constructed to ensure that the transformed series attains the target covariance $\Sigma^{counter}$. Further methodological details are provided in Appendix A.1. This transformation contracts the original series of realized and expected short-term inflation observations uniformly around the new mean, thereby preserving the empirical structure of co-movements across forecast horizons. The counterfactual series for realized inflation and expectations are generated by applying this transformation exclusively to post–Q1:2015 data, while all pre-IT observations remain unchanged.

Figure 6 presents the resulting counterfactual and actual inflation series.

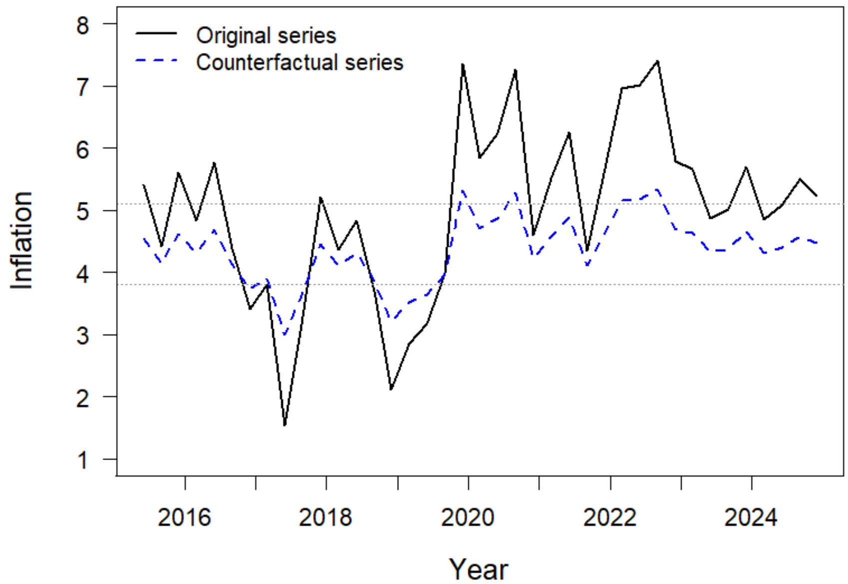

Figure 6: **Counterfactual and observed historical inflation series :**Counterfactual inflation path for the inflation-targeting period (post Q1:2015), constructed to reflect a lower-volatility regime aligned with the near-rational bound (3.8–5.1%). The observed inflation series is included for comparison.

Introducing this estimated counterfactual series into the dynamic learning algorithm produces the gain function sequence shown in Figure 7. The gain function under the counterfactual remains broadly similar to the benchmark during the pre-IT period; however, post-IT skewness declines substantially. This reduction occurs because inflation realizations remain consistently closer to the rational bound associated with the 4% policy target, thereby decreasing the frequency of rational-bound breaches that would otherwise increase long-run inflation uncertainty and induce asymmetric sensitivity in expectation formations.

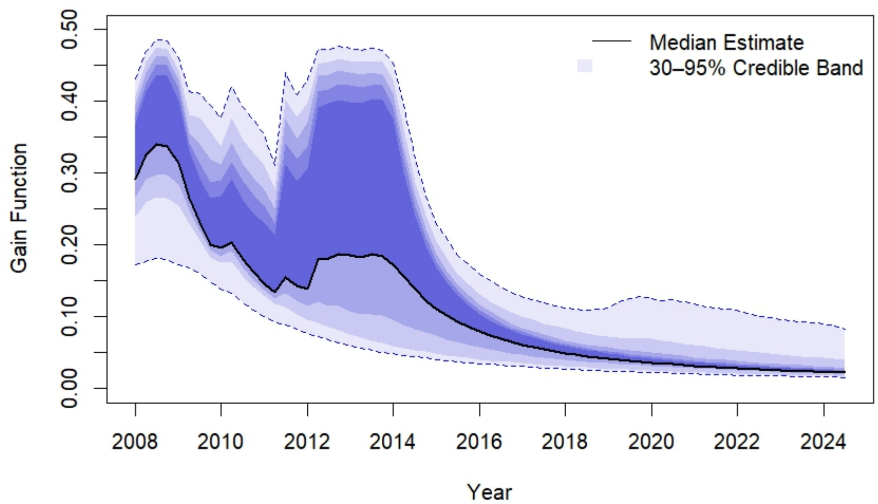

Figure 7: **Gain function implied by the counterfactual inflation series :**For the inflation-targeting period (post Q1:2015), the gain function is derived using the counterfactual inflation series; for earlier periods (pre Q2:2015), the observed historical inflation series is used. The shaded areas indicate the posterior interval from the 30th to the 95th percentile.

Next, a second counterfactual experiment is presented to capture the complementary mechanism influencing the sharp decline in long-run inflation expectations immediately after the adoption of IT. This mechanism is linked to the dynamics of the underlying inflation process itself. As shown in Figure 5, the transition to IT occurred during a period of subdued non-core inflation and a cyclical slowdown, as indicated by persistent declines in the OECD CLI . These conditions produced a sequence of negative inflation surprises, with realized inflation breaching the lower rational bound in few instances. The gain function in the immediate post-IT period remained moderately elevated due to agents' prior experience with significant positive inflation surprises. Consequently, the negative surprises that followed IT adoption exerted a strong downward effect on the perceived long-run inflation target through the update rule in equation (2.10). This mechanism amplified the downward drift in inflation target expectations within the PLM, reinforcing lower inflation expectations across the forecast horizon following the change in the inflation target. Collectively, these dynamics indicate that the sharp decline in inflation expectations, as widely noted in the Indian literature, does not necessarily indicate a fully anchored regime. Instead, such a decline can arise naturally under an expectations-formation process that remains only partially anchored.

To further substantiate this interpretation, a counterfactual experiment is conducted that imposes a monotonically declining gain function beginning in the first quarter of 2015. This approach effectively places agents directly into a fully anchored, decreasing-gain regime while maintaining all other estimated parameters and shock realizations identical to those in the benchmark scenario. The posterior band around the median gain estimates (Figure 8) mimics an environment in which all agents in the economy immediately after the IT transition adopt a completely forward-looking strategy and become immune to transient short-term inflation surprises within their rational bounds. The gain function becomes highly symmetric, and the posterior band becomes tight, reflecting low long-run inflation uncertainty. In such a scenario, the revisions in inflation target deviations becomes minimal as the gain function decreases monotonically. These lower innovations in the inflation target are factored into the agents' law of motion during the formation of long-term beliefs, resulting in a noticeably smaller decline in median long-run expectations between 2016 and 2020 than in the benchmark estimates. The sharp decline in the ten-year-ahead inflation expectations is underestimated by a huge margin.

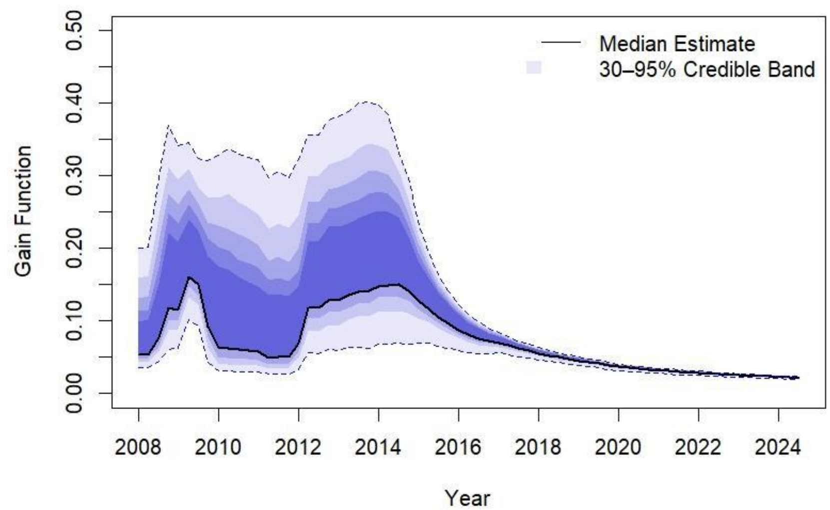

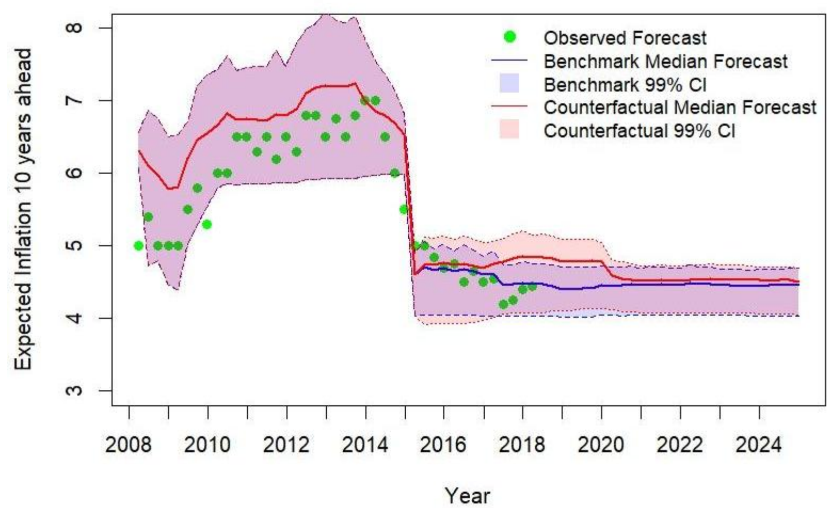

Figure 8: **Counterfactual gain function and implied long-horizon inflation expectations under a monotonically declining gain regime post Q1:2015 :**The upper panel shows the counterfactual gain function. The lower panel displays the counterfactual median estimate of ten-year-ahead inflation expectations and the associated 99% posterior band, compared to the benchmark estimates.

In summary, the findings suggest that forecasters in India remain in an adaptive-learning phase, even in the the current inflation-targeting regime. Inflation expectations across the forecast horizon continue to adjust in response to both positive and negative inflation surprises. These breaches, which primarily occur on the upside and are disproportionately concentrated during periods of adverse growth and cyclical slowdowns, perpetuate long-horizon uncertainty and impede a rapid transition to a fully anchored expectations regime. As a result, the gain function does not converge to its limiting, low-volatility level but instead maintains sufficient skewness to transmit short-term inflation surprises, relative to the perceived rational bounds, into revisions of the long-run expectations.

4.3. Bond Market Implications

The observed inflation dynamics extend to the behavior of the term-premium component of bond yields. As established in the previous subsection, breaches of the rational bound occur primarily during economic downturns, when inflation overshoots coincide with weak contemporaneous output. During these periods, upward revisions to inflation expectations across horizons are typically negatively correlated with consumption growth. These repeated negative covariances generate long-run risk for bondholders, which is reflected in the cross-section of bond yields as a persistent term premium, particularly at longer maturities. Such long-run risks mirror the uncertainty captured by the upward skew in the distribution of the gain function. Asymmetric skewness in the gain function, which reflects frequent upward revisions in inflation beliefs, directly translates into the persistence of term premia across maturities.

Appendix B formalizes this mechanism within a forward-looking long-run-risk asset-pricing framework, as outlined by Song (2017) and Bansal and Yaron (2004). In both pre- and post-inflation-targeting (IT) regimes, replicating the observed term premium requires incorporating persistent updates in agents' perceived inflation target, driven by innovations that covary with trends in real consumption growth. Estimates from the forward-looking model indicate that the central bank, even under an explicit inflation-targeting regime, does not actively stabilize inflation. In a counterfactual scenario where the macroeconomic environment remains constant and only the degree of inflation stabilization reflects active stabilization, the covariance channel weakens significantly, and the ten-year term premium decreases by approximately 120 basis points. Thus, the persistence of the term premium, especially at longer maturities, is directly attributable to inflation uncertainty stemming from continual innovations in inflation target expectations. Reducing such evolution through active inflation stabilization proportionally compresses the term premium.

5. Conclusion

This paper addresses three issues: how long-term inflation expectations in India have evolved following the adoption of inflation targeting; the extent to which long-run beliefs are anchored to the 4 percent policy target; and whether uncertainty in these beliefs contributes to the persistence of term premia, particularly at longer maturities of government securities. The results show that ten-year-ahead inflation expectations have declined and moved closer to the policy target after the introduction of inflation targeting. Revisions to long-run expectations have also fallen substantially relative to the pre-inflation-targeting period, when the absence of an explicit target left expectations weakly anchored.

Despite these improvements, long-run inflation expectations remain imperfectly anchored. Some agents remain only partially forward-looking, and long-term beliefs continue to respond to short-run inflation surprises. This persistence reflects repeated breaches of the rational bounds agents place around their perceived inflation target. For a 4 percent target, these bounds are estimated to lie between 3.8 and 4.9 percent, narrower than the central bank's official tolerance band. Greater stabilization of inflation around agents' perceived target reduces long-run inflation uncertainty, lowers the term premium, and strengthens the anchoring of long-term expectations.

Appendix A Adaptive Learning Model Summary

The reduced system of equations stated in Section 2 can be re-written as:

$$\hat{\pi}_t = (1-\gamma)\Gamma\big[1 + g_t^{-1}(1-\gamma)(\Gamma-1)\big]\,\bar{\pi}_{t-1} + \gamma_p\hat{\pi}_{t-1} + \rho_\epsilon\epsilon_{t-1} + \eta_{t-1},$$

$$\bar{\pi}_t = \bar{\pi}_{t-1}\big[1 + g_t^{-1}(1-\gamma)(\Gamma-1)\big] + g_t^{-1}(\varepsilon_{t-1} + \mu_{t-1}),$$

$$\eta_t = \mu_t + \varepsilon_t,$$

$$\epsilon_t = \rho_\epsilon\epsilon_{t-1} + \varepsilon_t.$$

And, the gain evolves according to:

$$g_t = \begin{cases} g_{t-1} + 1, & \Big|\hat{E}_{t-1}\hat{\pi}_t - E_{t-1}\hat{\pi}_t\Big| \leq \Theta\sqrt{\varepsilon_t^2 + \mu_t^2} \quad \textbf{(Decreasing gain)}, \\ \bar{g}^{-1}, & \textbf{otherwise (Constant gain)}. \end{cases}$$

The system contains linear state variables $\varsigma_t = (\hat{\pi}_t, \eta_t, \epsilon_t)$ and nonlinear states $\bar{\pi}_t$ and g_t. The linear state-space is ultimately represented as:

$$\varsigma_t' \;=\; A \;+\; B\,\varsigma_{t-1}' \;+\; S\,(\varepsilon_t, \mu_t)', \tag{A.1}$$

with

$$A = \begin{bmatrix} \bar{\pi}_{t-1}\big[1 + g_t^{-1}(1-\gamma)(\Gamma-1)\big] \\ 0 \\ 0 \end{bmatrix}, \qquad B = \begin{bmatrix} \gamma_p & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \rho_\epsilon \end{bmatrix}, \qquad S = \begin{bmatrix} 0 & 0 \\ 1 & 1 \\ 1 & 0 \end{bmatrix}.$$

It can be seen the linear state-space system are conditional on the non-linear state variables. The nonlinear state variables are initially marginalized out, and approximated using the particle filters wherein each particle evolve as:

$$\begin{aligned}\bar{\pi}_t &= \bar{\pi}_{t-1}\Big[1 + \big(I_{t-1}\bar{g}^{-1} + (1 - I_{t-1})(g_{t-1} + 1)\big)^{-1}(1-\gamma)(\Gamma - 1)\Big] \\ &\quad + \big(I_{t-1}\bar{g}^{-1} + (1 - I_{t-1})(g_{t-1} + 1)\big)^{-1}\ e_2'(\varsigma_{t-1}), \qquad \text{(A.2)} \\ g_t &= I_{t-1}\bar{g}^{-1} + (1 - I_{t-1})(g_{t-1} + 1).\end{aligned}$$

Conditional on each updated particle, the linear state variables in the equation (A1) are updated using the Kalman filter.

The mixed linear and nonlinear state space system are linked to realized inflation and one- to four-quarter-ahead SPF expectations as observables. The measurement matrix in the equation (2.13) for the nonlinear states $\bar{\pi}_t$ and g_t is

$$C_{1,t} = \begin{bmatrix} (1-\gamma)\Gamma & (1-\gamma^2)\Gamma & (1-\gamma^3)\Gamma & (1-\gamma^4)\Gamma & (1-\gamma^5)\Gamma \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

and for the linear states ς_t:

$$C_{2,t} = \begin{bmatrix} \gamma & 0 & \rho \\ \gamma^2 & 0 & \rho^2 + \gamma\rho \\ \gamma^3 & 0 & \rho^3 + \gamma\rho^2 + \gamma^2\rho \\ \gamma^4 & 0 & \rho^4 + \gamma\rho^3 + \gamma^2\rho^2 + \gamma^3\rho \\ \gamma^5 & 0 & \rho^5 + \gamma\rho^4 + \gamma^2\rho^3 + \gamma^3\rho^2 + \gamma^4\rho \end{bmatrix}.$$

The coefficient matrix corresponding to the measurement errors follow $D_t = I_5$.

A.1. Counterfactual Inflation Series

A.1.1. Estimation of matrix ∇

The sample standard deviations of the observed series of the realized and short-term inflation expectations can be denoted by, $U = \mathrm{diag}(\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5)$, and the corresponding empirical correlation matrix is, $R = U^{-1}\Sigma U^{-1}$. To generate the counterfactual series this empirical correlation matrix is first shrinked by a factor of ω=0.4, while preserving the empirical correlation structure $U^{counter} = \omega U,\ \Sigma^{counter} = U^{counter} R U^{counter}$.

Cholesky decomposition is used to transform the empirical covariance $\Sigma = LL^{\top},\ \Sigma^{counter} = L^{counter} L^{counter\top}$.

Then the corresponding linear transformation matrix that maps the empirical covariance Σ into the shrunken covariance $\Sigma^{counter}$ is $\nabla = L^{counter} L^{-1}$.

A.1.2. Construction of Target Means and Counterfactual Inflation Series

Let the empirical post-Q1:2015 means be represented as: $\mu = \begin{pmatrix} \mu_{spot} & \mu_1^{SPF} & \mu_2^{SPF} & \mu_3^{SPF} & \mu_4^{SPF} \end{pmatrix}$. wherein the empirical gaps between realized inflation and one to four quarter ahead inflation expectation is defined as: $\wp_h = \mu_{spot} - \mu_h^{SPF}, h = 1, \ldots, 4.$

The counterfactual mean of realized inflation is set at 4.45%, the mid point of the near rational bound of 3.8% to 5.1% , i.e., $\mu_{spot}^{counter} = \bar{\pi} = 4.45$, and the counterfactual one to four quarter ahead mean expectations are also adjusted as per the empirical gaps: $\mu_h^{SPF,counter} = \bar{\pi} - \wp_h$. Post Q1:2015, the affine transformation of original series Y_t gives counterfactual inflation series: $Y_t^{\text{counter}} = \mu^{counter} + \nabla(Y_t - \mu^{counter}), t >$

t_{2015M3}.

A.1.3. Plots of Food Inflation and Oil Price around the implementation of IT regime

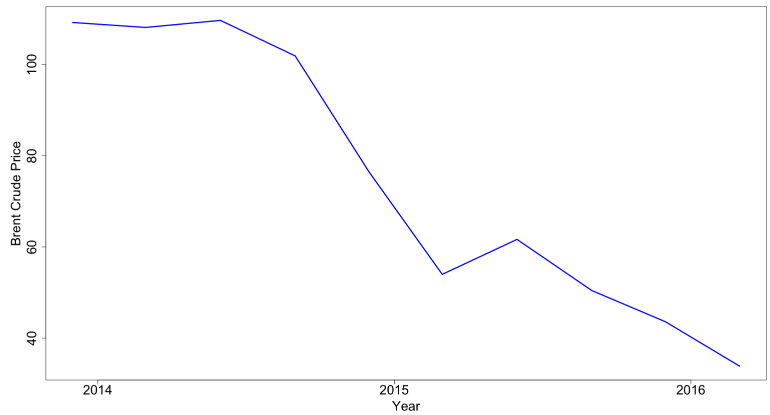

Figure A1: The figure displays the Brent Crude Spot Prices for the period 2014–2016, covering the entire transition period to the inflation targeting regime.

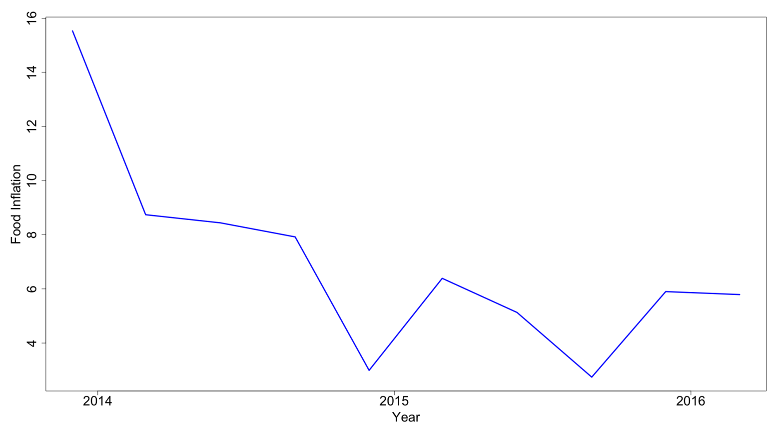

Figure A2: The figure displays the Food Inflation of India for the period 2014–2016. The food CPI index is sourced from the Ministry of Statistics and Programme Implementation.

Appendix B Asset Pricing Model

This asset-pricing model integrates the real long-run risk framework of Bansal and Yaron (2004) and the nominal long-run risk from Song (2017) within an endowment economy. The model incorporates a representative agent with recursive preferences, as in Epstein and Zin (1989) and Weil (1989). In this framework, risk is priced ex ante by the representative agent when forming preferences, primarily through early resolution of innovations in two latent states: drifts in perceived inflation targets (nominal side) and trends in real consumption growth (real side). These latent states drive the stochastic discount factor in the framework, thereby generating nominal and real risk premia.

In contrast to the framework presented in the main text, this model is purely forward-looking and rational. However, by incorporating volatility into the latent states and employing a Bayesian framework, the model captures a component of real-time uncertainty, which can be interpreted as the economy's inherent long-term feature, as captured by the adaptive learning model. In our case, through the learning framework, agents learn the recurrent counter cyclical innovations in the inflation-target beliefs as the long-term feature of the underlying economy. This asset pricing framework examines the extent to which this long-run feature accounts for the observed ten-year term premium across both pre- and post-inflation targeting (IT) regimes.

B.1. The Framework and the latent states

B.1.1. Framework and Latent States

Under Epstein–Zin–Weil preferences, the expected return $R_{i,t+1}$ on any observable asset satisfies the Euler equation

$$E_t\Big[\delta^\theta (\Delta C)_{t+1}^{-\theta/\Psi}\, R_{c,t+1}^{-(1-\theta)}\, R_{i,t+1}\Big] = 1, \tag{B.1}$$

where $(\Delta C)_{t+1}$ denotes aggregate real consumption growth and $R_{c,t+1}$ is the gross real return on the (unobservable) consumption-claim asset. The parameters are: δ (time-discount factor), Ψ (intertemporal elasticity of substitution), and $\theta = (1-\gamma)/(1-\Psi)$, with γ the coefficient of relative risk aversion. The expected returns via SDF are linked to two latent states- a) Inflation Target perceptions, and b) Trends in real consumption growth. Innovations to these states drive movements in asset returns and premiums.

The reduced-form system of equations defining the latent states are as follows:

$$\begin{aligned}
G_{t+1} &= [\delta c_t,\ \delta d_t]^\top, \\
G_{t+1} &= \mu + \varphi Y_t + \Sigma \eta_{t+1}, \qquad \eta_{t+1} \sim \mathcal{N}(0, I), \\
Y_{t+1} &= \Psi_1 Y_t + \Psi_2 \Sigma_y \xi_{y,t+1}, \\
Y_t &= [y_{c,t},\ y_{\pi,t}]^\top, \\
\xi_{y,t+1} &= [\xi_{yc,t},\ \xi_{y\pi,t}]^\top.
\end{aligned} \tag{B.2}$$

$$\mu = \begin{pmatrix} \mu_c & \mu_d \end{pmatrix}, \qquad \varphi = \begin{pmatrix} 1 & 0 \\ \phi & 0 \end{pmatrix}, \qquad \Psi_1 = \begin{pmatrix} \varphi_c & 0 \\ 0 & \varphi_\pi \end{pmatrix}, \qquad \Psi_2 = \begin{pmatrix} 1 & 0 \\ \alpha & 1 \end{pmatrix},$$

$$\Sigma_y = \begin{pmatrix} \sigma_{yc} & 0 \\ 0 & \sigma_{y\pi} \end{pmatrix}.$$

The latent states $y_{c,t}$ and $y_{\pi,t}$ represent the latent component of real consumption growth and the drift in the perceived inflation target. Particularly the innovation in $y_{\pi,t}$ component is of interest to us, and the term α determines the cyclicality of innovations in the inflation-target with respect to the innovations in real consumption trend. In the setup δc_t and δd_t denote real growth in aggregate consumption and dividends. The dividend component loads on consumption growth with leverage ϕ, capturing aggregate equity-market evidence. Parameters μ_c and μ_d are average growth rates in real consumption and dividends.

B.1.2. Monetary Policy Rule and Nominal Interest Rates

Monetary policy follows a Taylor-type rule, with the central bank adjusting the short-term nominal interest rate in response to trends in real consumption growth and deviations of inflation from its perceived target. The nominal rate i_t is specified as

$$i_t = \chi_0 + \chi_c y_{c,t} + \chi_\pi \big(\pi_t - \Theta_0 - y_{\pi,t}\big) + y_{\pi,t} + y_{i,t}, \tag{B.3}$$

where $y_{i,t}$ is an exogenous Taylor-rule shock following an AR(1) process. The coefficient χ_π reflects the strength of inflation stabilization around the perceived target $y_{\pi,t}$. When policy is weakly stabilizing—especially amid countercyclical inflation-target drifts ($\alpha < 0$)—agents demand a premium, wherein the nominal component of premium is driven by the degree of policy deviation from active stabilization and magnitude of shocks to the perceived inflation-target.

Bond prices are modeled in affine form with respect to the latent states, following Song (2017) and Doh (2013). The imposed no-arbitrage condition across maturities generates the underline term structure. Innovations in the latent states affect the SDF, influencing both bond prices and the term premium.

B.1.3. Data, Parameter Estimation and Discussion

Real consumption growth is calculated as the four-quarter log change in Real Private Final Consumption Expenditure, with data obtained from the CEIC database. Zero-coupon yield curve (ZCYC) yields are derived using Nelson–Siegel–Svensson parameters provided by the Clearing Corporation of India Limited. Inflation is measured as the four-quarter logarithmic change in the Consumer Price Index (CPI). Equity returns and price-to-dividend ratios are computed using the NSE 500 index, which represents approximately 90% of free-float market capitalization. The observation vector $Y_t = (\Delta c_t, \pi_t, pd_t, y_{1,t}, y_{2,t}, y_{5,t}, y_{10,t})$ includes consumption growth, inflation, price-to-dividend ratio, and ZCYC yields at five maturities. The model is estimated at a quarterly frequency from Q3:2005 to Q4:2023 using adaptive Bayesian methods, achieving convergence at a 25.4% acceptance rate over 200,000 iterations.

Table A1: **Priors and Posterior Parameter Distributions of the Asset-Pricing Model**

	Prior			Posterior		
Parameter	Dist.	1^{th} perc	99^{th}perc	Mode	1^{th} perc	99^{th} perc
Estimated parameters (Bayesian update)						
α	U	-5	-0.6	-0.619	-0.72	-0.6
χ_0	G	0.059	0.07	0.067	0.062	0.069
χ_π	G	0.75	1.2	0.916	0.915	0.917
χ_c	G	0.17	0.5	0.337	0.33	0.34
σ_{yc}	IG	0.0015	0.0019	0.00189	0.0018	0.0019
$\sigma_{y\pi}$	IG	0.0007	0.0019	0.00189	0.00018	0.0019
σ_{yi}	IG	0.0007	0.00019	0.00188	0.0018	0.0019
Fixed parameters (not updated in Bayesian step)						
φ_c	fixed	value: *0.99*			—	
φ_π	fixed	value: *0.995*			—	
φ_i	fixed	value: *0.88*			—	
γ	fixed	value: *8*			—	
ψ	fixed	value: *2*			—	
μ_c	fixed	value: *0.06*			—	
μ_d	fixed	value: *0.052*			—	
ϕ	fixed	value: *0.85*			—	

Notes: U = Uniform; G = Gaussian; IG = Inverse-Gamma. The measurement errors are fixed at 0.8% of the observed sample std error. μ_c and μ_d are fixed to sample average. ϕ is fixed to the ratio of the standard deviation observed in the sample for real dividend growth to real consumption growth.The measurement errors are fixed at 0.8% of the standard deviation of the observables.

Table A2: Model vs. Data: Quarterly Asset-Pricing Moments; Full Sample (Q3:2005 to Q4:2023) and Post-IT (Q1:2015 to Q4:2023)

	Data				Model			
	Full sample		Post-IT		Full sample		Post-IT	
Variable	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
$y_{1,t}$	6.53	1.4	5.66	1.2	6.51	1.19	5.73	0.96
$y_{2,t}$	6.78	1.2	5.87	1.1	6.77	1.11	6.06	0.90
$y_{5,t}$	7.23	0.84	6.18	0.75	7.24	0.91	6.72	0.75
$y_{10,t}$	7.59	0.69	6.78	0.50	7.57	0.72	7.28	0.60
r_m	15.29	—			14.33	—		
$\sigma(r_m)$	0.24	—			0.20	—		
ACF$(y_{10,t})$	0.68	—			0.80	—		
Term premium (bps)	145	—	136		125	—	125	

Notes: $y_{1,t}$, $y_{2,t}$, $y_{5,t}$, $y_{10,t}$ are the ZCYC yields of one, two, five and ten years maturity respectively. r_m is the aggregate equity market return. Term premium estimates are in basis points.

Table A3: Counterfactual Asset-Pricing Moments

Variable	Data	Case I	Case II
$y_{1,t}$	6.53	6.49	6.51
$y_{2,t}$	6.78	6.71	6.77
$y_{5,t}$	7.23	7.21	7.26
$y_{10,t}$	7.59	7.61	7.56
r_m	15.29	14.33	14.00
$\sigma(r_m)$	0.24	0.20	0.20
Term premium (bps)	145	-113	97

Notes: Case I sets inflation stabilization parameter in the inflation path to $\xi_p = 1.05$, holding all other parameters at posterior means in the benchmark estimates. Case II reduces α by 50%, again holding all other parameters at posterior means.

The model accurately reflects key asset-pricing moments. Countercyclical drifts in perceived inflation target, alongwith policy geared away from active inflation stabilization generate a persistent term premium, aligning with sample averages. Active inflation stabilization, which is analogous to an environment of anchored expectations, reduces the term premium by about 120 basis points, primarily by dampening the negative covariance channel. These results, detailed in Tables A2 and A3, are consistent across both pre- and post-IT regimes.

References

Adrian, T., Crump, R. K., and Moench, E. (2015). Regression-based estimation of dynamic asset pricing models. *Journal of Financial Economics*, 118(2):211–244.

Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society Series B*, 72(3):269–342.

Ascari, G., Phaneuf, L., and Sims, E. (2016). Business Cycles, Investment Shocks, and the "Barro-King" Curse. NBER Working Papers 22941, National Bureau of Economic Research, Inc.

Asnani, S., Kumar, P., and Tomar, S. (2019). Does Inflation Targeting Anchor Inflation Expectations? Evidence from India. *Working Paper. The Indian Statistical Institute.*

Ball, L. and Mankiw, N. G. (1995). Relative-price changes as aggregate supply shocks. *The Quarterly Journal of Economics*, 110(1):161–193.

Bansal, R. and Shaliastovich, I. (2013). A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets. *The Review of Financial Studies*, 26(1):1–33.

Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509.

Bernanke, B. S. (2007). Inflation expectations and inflation forecasting. Speech 306, Board of Governors of the Federal Reserve System (U.S.).

Branch, W. and Evans, G. (2011). Monetary policy and heterogeneous expectations. *Economic Theory*, 47(2):365–393.

Cai, 2010). A metropolis-hastings robbins-monro algorithm for maximum likelihood nonlinear latent structure analysis with a comprehensive measurement model. *Psychometrika*, 75(4):613–637.

Carter, C. K. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Carvalho, C., Eusepi, S., Moench, E., and Preston, B. (2023). Anchored inflation expectations. *American Economic Journal: Macroeconomics*, 15(1):1–47.

Chakraborty, A. and Evans, G. W. (2008). Can perpetual learning explain the forward-premium puzzle? *Journal of Monetary Economics*, 55(3):477–490.

Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018). The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature*, 56(4):1447–91.

Doh, T. (2013). Long-run risks in the term structure of interest rates: estimation. *Journal of Applied Econometrics*, 28(3):478–497.

Doucet, A., de Freitas, N., and Gordon, N. (2001). *An Introduction to Sequential Monte Carlo Methods*, pages 3–14. Springer New York, New York, NY.

Eichengreen, B., Gupta, P., and Choudhary, R. (2021). Inflation Targeting in India: An Interim Assessment. *India Policy Forum*, 17(1):77–141.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica: Journal of the Econometric Society*, pages 937–969.

Eusepi, S. and Preston, B. (2010). Central bank communication and expectations stabilization. *American Economic Journal: Macroeconomics*, 2(3):235–71.

Evans, G. W. and Honkapohja, S. (2003). Adaptive learning and monetary policy design. *Journal of Money, Credit and Banking*, pages 1045–1072.

Garga, V., Lakdawala, A., and Sengupta, R. (2022). Assessing central bank commitment to inflation targeting: Evidence from financial market expectations in India. Indira Gandhi Institute of Development Research, Mumbai Working Papers 2022-017, Indira Gandhi Institute of Development Research, Mumbai, India.

Guse, E. A. (2014). Adaptive learning, endogenous uncertainty, and asymmetric dynamics. *Journal of Economic Dynamics and Control*, 40:355–373.

Honkapohja, S. and Mitra, K. (2003). Learning with bounded memory in stochastic models. *Journal of Economic Dynamics and Control*, 27(8):1437–1457.

Kim, Y. S. (2009). Exchange rates and fundamentals under adaptive learning. *Journal of Economic Dynamics and Control*, 33(4):843–863.

Kishor, N. K. and Pratap, B. (2023). The role of inflation targeting in anchoring long-run inflation expectations: Evidence from india. *Available at SSRN 4633861*.

Kitagawa, G. (1993). A monte carlo filtering and smoothing method for non-gaussian nonlinear state space models. In *Proceedings of the 2nd US-Japan joint seminar on statistical time series analysis*, volume 110.

Lansing, K. (2009). Time varying u.s. inflation dynamics and the new keynesian phillips curve. *Review of Economic Dynamics*, 12(2):304–326.

LeBaron, B. and Smith, K. (2025). Learning integrated inflation forecasts in a simple multi-agent macroeconomic model. *Journal of Economic Dynamics and Control*, 172(C).

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

Milani, F. (2007). Expectations, learning and macroeconomic persistence. *Journal of monetary Economics*, 54(7):2065–2082.

Pattanaik, S., Nadhanael, G., and Muduli, S. (2023). Taming Inflation by Anchoring Inflation Expectations. *Economic and Political Weekly*, Vol. 58(Issue No. 22).

Preston, B. (2008). Adaptive learning and the use of forecasts in monetary policy. *Journal of Economic Dynamics and Control*, 32(11):3661–3681.

Rotemberg, J. J. (1982). Sticky Prices in the United States. *Journal of Political Economy*, 90(6):1187–1211.

Schön, T. B., Gustafsson, F. K., and Nordlund, P.-J. (2005). Marginalized particle filters for mixed linear/nonlinear state-space models. *IEEE Transactions on Signal Processing*, 53:2279–2289.

Song, D. (2017). Bond Market Exposures to Macroeconomic and Monetary Policy Risks. *The Review of Financial Studies*, 30(8):2761–2817.

Stock, J. H. and Watson, M. W. (2007). Why has u.s. inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39(s1):3–33.

Vihola, M. (2012). Robust adaptive metropolis algorithm with coerced acceptance rate. *Statistics and Computing*, 22(5):997–1008.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of monetary economics*, 24(3):401–421.

Woodford, M. (2003). Optimal interest-rate smoothing. *The Review of Economic Studies*, 70(4):861–886.

Woodford, M. and Walsh, C. E. (2005). Interest and prices: Foundations of a theory of monetary policy. *Macroeconomic Dynamics*, 9(3):462–468.