Contextual Expectations and Inflation Dynamics

Ceyhun Elgin*

Bogazici University

Abstract

This paper develops a Contextual Expectations (CE) framework in which agents form inflation expectations by combining multiple signals with endogenous, context-dependent weights governed by a contextual relevance function (trustworthiness, relevance, prominence, similarity to recent experience) and a memory parameter that generates persistence. CE nests adaptive and rational expectations as limiting cases while allowing time-varying emphasis on backward-looking versus policy-driven information. Embedded in a New Keynesian inflation setting, CE implies that changes in the weight on central bank communication shape inflation persistence, re-anchoring dynamics, and monetary policy effectiveness. We calibrate the model to Turkey (2013-2023) and illustrate via simulations how credibility and target shocks propagate through endogenous weight adjustment. A randomized information-update experiment provides causal evidence on how central bank guidance, backward inflation information, and price-salience news shift expectations and vary with contextual dimensions.

Keywords: expectation formation; monetary policy; inflation; survey experiment

JEL codes: D84; E31; E52

1 Introduction

Expectation formation is a cornerstone of economic decision-making, playing a crucial role in shaping macroeconomic outcomes. The way economic agents form their expectations about future inflation rates has crucial implications for consumption, investment, and policy effectiveness. Traditionally, macroeconomic models have relied on simplified assumptions about expectation formation, with rational expectations (RE) and adaptive expectations (AE) being the two dominant paradigms. The RE hypothesis, pioneered by Muth (1961)

* Address: Bogazici University, Department of Economics, Natuk Birkan Kat 2 Bebek, Istanbul, Turkey. E-Mail: ceyhunelgin@gmail.com

and later popularized by Lucas (1972), posits that economic agents use all available information efficiently and have model-consistent views of the economy. This approach assumes that individuals' predictions are, on average, correct and that systematic errors in forecasting are quickly eliminated as agents learn from their mistakes. While powerful in its simplicity, the rational expectations hypothesis has faced criticism for its strong assumptions about individuals' cognitive abilities and information processing capabilities (Evans and Honkapohja, 2001). On the other hand, the AE model, introduced by Cagan (1956) and further developed by Friedman (1957), suggests that individuals form their expectations based on past experiences, gradually adjusting their predictions as new information becomes available. This approach captures the learning process more realistically but has been criticized for its potential to lead to systematic forecasting errors, especially in rapidly changing economic environments (Sargent, 1993).

However, both these models may overlook the complexity and dynamic nature of real-world expectation formation. They fail to account for the complex context in which economic decisions are made, including varying levels of trust in information sources, personal financial situations, and overall economic perceptions. This gap in existing models calls for a more subtle approach to understanding how inflation expectations are formed. For this purpose, this paper introduces a novel theoretical framework called Contextual Expectations (CE), which aims to integrate multiple dimensions of information and individual context to better understand the formation of inflation expectations. By incorporating economic perceptions, trust in information sources, personal financial situations, and outlook on future economic conditions, contextual expectations provide a more comprehensive and realistic depiction of expectation formation.

While sticky information models (Mankiw and Reis, 2002) rely on the infrequent updating of information sets, and rational inattention models (Sims, 2003) emphasize capacity constraints in processing available data, the CE framework focuses on the *weighting* of information based on its perceived context. In our model, agents may well see the central bank's signal (no information stickiness) and have the capacity to process it (no inattention), yet still choose to underweight it if the context; defined by trust, relevance, and similarity, renders that signal less compelling than backward-looking alternatives. This distinction allows

CE to explain why information transmission varies not just with frequency or cost, but with the evolving credibility of the source.

The proposed contextual expectations model builds upon insights from behavioral economics, particularly the work of Kahneman and Tversky (1979) on prospect theory, which emphasizes the importance of reference points and framing in decision-making. It also draws inspiration from Simon's (1955) concept of bounded rationality, acknowledging the cognitive limitations of economic agents and the importance of satisficing behavior in complex environments. By developing this theoretical framework, we aim to bridge the gap between the simplifying assumptions of traditional models and the complex reality of how individuals form their expectations about future inflation. This approach has the potential to enhance our understanding of inflation dynamics, improve the effectiveness of monetary policy, and contribute to more accurate macroeconomic modeling.

The remainder of this paper is structured as follows: Section 2 provides a short literature review on expectation formation models. Section 3 presents the theoretical framework of contextual expectations, including its mathematical formulation and the concept of the contextual relevance function; it also derives implications for inflation dynamics, monetary policy, and macroeconomic stability, and situates CE relative to rational and adaptive benchmarks. Section 4 documents motivating stylized facts and reports the calibration and simulation analysis. Section 5 presents a randomized information-update survey experiment that identifies the causal effects of central bank guidance, backward-looking inflation information, and price-salience news on twelve-month-ahead inflation expectations, explores heterogeneity in updating along the model's contextual dimensions, and maps the experimental treatment effects into empirical counterparts of the CE signal weights. Finally, Section 6 discusses the advantages and challenges of the contextual expectations approach, followed by some concluding remarks.

2 Literature Review

The study of expectation formation in economics has a rich history, evolving significantly over the past century. This literature review traces the development of key expectation formation models, focusing on their assumptions, strengths, and criticisms, as well as recent advancements in the field. The earliest formal treatment of expectations in economic theory can be traced back to Keynes (1936), who emphasized the role of psychological factors and uncertainty in shaping economic decisions. Keynes introduced the concept of "animal spirits" to describe the spontaneous urge to action that drives economic behavior, highlighting the importance of expectations in determining investment and consumption decisions. Building on Keynes' insights, the adaptive expectations hypothesis gained prominence in the 1950s and 1960s. Cagan (1956) introduced this concept in his study of hyperinflation, proposing that individuals form their expectations of future inflation based on past inflation rates. The adaptive expectations model was further developed by Nerlove (1958) and applied to various economic contexts, including Friedman's (1957) permanent income hypothesis. The adaptive expectations model assumes that economic agents revise their expectations gradually as new information becomes available. Mathematically, it can be expressed as:

$$E[\pi_{t+1}] = \lambda \pi_t + (1 - \lambda) E[\pi_t], \tag{1}$$

where $0 < \lambda < 1$.

Here, $E[\pi_{t+1}]$ represents the expected inflation rate for the next period, π_t is the current period's inflation rate, and λ is the adjustment parameter. This formulation captures the idea that individuals learn from their past prediction errors, albeit slowly. While the adaptive expectations model provided a tractable framework for modeling expectations, it faced criticism for its potential to lead to systematic forecasting errors, especially in rapidly changing economic environments. Critics argued that it failed to account for forward-looking behavior and the use of all available information by economic agents. In response to these limitations, Muth (1961) introduced the rational expectations hypothesis, which was later popularized and extended by Lucas (1972) and Sargent and Wallace (1975). The rational expectations theory posits that economic agents use all available information efficiently and have model-consistent views of the economy. Under this hypothesis, the expected value of inflation conditional on the agent's information set equals the mathematical expectation given all available information:

$$E_t[\pi_{t+1}] \equiv E[\pi_{t+1} \mid \Omega_t] \tag{2}$$

where Ω_t denotes the full information set available at time t, encompassing all relevant economic variables, policy announcements, and historical data. This formulation implies that, on average, economic agents' expectations are correct, and any forecast errors are orthogonal to the information set.

The rational expectations hypothesis had a profound impact on macroeconomic theory and policy analysis. It led to the development of the Lucas critique (Lucas, 1976), which argued that traditional econometric models were inadequate for policy analysis because they failed to account for the way policy changes alter the structure of economic relationships. Despite its theoretical elegance, the rational expectations hypothesis has faced considerable criticism. Simon (1955) introduced the concept of bounded rationality, arguing that individuals have limited cognitive abilities and often settle for satisfactory rather than optimal solutions. Kahneman and Tversky's (1979) prospect theory further challenged the assumptions of rational expectations by demonstrating systematic biases in human decision-making under uncertainty. In response to these criticisms, several alternative approaches to modeling expectations have emerged. Evans and Honkapohja (2001) developed models of adaptive learning, where agents behave like econometricians, continuously updating their forecasting models as new data becomes available. This approach bridges the gap between adaptive and rational expectations, allowing for a more realistic depiction of the learning process. Sargent (1993) introduced the concept of bounded rationality in macroeconomics, proposing models where agents have limited computational capabilities and use simple forecasting rules. This approach has been further developed by researchers like Hommes (2011), who studied heterogeneous expectations models where different agents use various forecasting strategies. More recently, Gabaix (2020) developed a behavioral New Keynesian model in which agents exhibit cognitive discounting—they underweight future variables due to bounded rationality, generating a "cognitive discount factor" that produces more realistic impulse responses than standard RE models. Angeletos and Huo (2021) decompose expectation dynamics into "myopia" (underreaction to future fundamentals) and "anchoring" (persistence tied to past beliefs), providing a framework that resonates with our distinction between backward-looking similarity and forward-looking trust in policy signals. Our CE model complements these approaches by endogenizing the weights agents place on different information sources as functions of contextual characteristics, rather than treating cognitive limitations as fixed parameters.

Recent research has also explored the role of information frictions in expectation formation. Mankiw and Reis (2002) proposed the sticky information model, where only a fraction of agents update their information set in each period. Similarly, Sims (2003) developed the rational inattention theory, arguing that individuals have limited capacity to process information and must choose which information to pay attention to. Empirical studies have provided mixed evidence on the nature of expectation formation. While some studies, such as Lovell (1986), found evidence supporting adaptive expectations, others, like Keane and Runkle (1990), found support for rational expectations in certain contexts. Many researchers, including Branch (2004), have found evidence of heterogeneity in expectation formation, with different individuals using different forecasting strategies.

More recently, the literature has focused on how economic agents perceive and react to specific communication strategies. Coibion, Gorodnichenko, and Weber (2022) highlight how households often disregard monetary policy announcements unless they directly relate to their personal economic context, a finding that resonates with our model's "relevance" parameter. Similarly, D'Acunto et al. (2021) demonstrate that diversity in cognitive abilities and demographic characteristics leads to substantial heterogeneity in how policy signals are interpreted, further motivating our focus on agent-specific contextual weights.

The ongoing debate and evolving research in expectation formation underscore the need for more nuanced models that can capture the complexity of real-world decision-making. The contextual expectations framework proposed in this paper aims to address this need by incorporating multiple dimensions of information and individual context into the expectation formation process.

3 Theoretical Framework

3.1 Foundations of Contextual Expectations

The contextual expectations framework builds upon existing literature while incorporating insights from behavioral economics and information theory. We formalize this approach mathematically, recognizing that economic agents form expectations within a complex, multidimensional context.

3.2 Model Specification

Let π_t denote the inflation rate at time t, and $E_t[\pi_{t+1}]$ the expectation of inflation for period $t+1$ formed at time t. We define the contextual expectations model as:

$$E_t[\pi_{t+1}] = \sum_{i=1}^{n} w_{i,t} \cdot I_{i,t} \tag{3}$$

where:

- $w_{i,t}$ is the weight assigned to information source i at time t
- $I_{i,t}$ is the inflation signal from information source i at time t
- n is the total number of information sources

3.3 Contextual Relevance Function

The weights $w_{i,t}$ are determined by the contextual relevance function, defined as:

$$w_{i,t} = \frac{C_{i,t}}{\sum_{j=1}^{n} C_{j,t}} \tag{4}$$

where $C_{i,t}$ is the contextual relevance of information source i at time t. By construction, $w_{i,t} \geq 0$ and $\sum_{i=1}^{n} w_{i,t} = 1$, so that the CE rule produces a convex combination of signals. We define $C_{i,t}$ as:

$$C_{i,t} = f(T_{i,t}, R_{i,t}, P_{i,t}, S_{i,t}) \tag{5}$$

where:

- $T_{i,t}$ represents the perceived trustworthiness of the information source
- $R_{i,t}$ represents the perceived relevance of the information to the current economic situation
- $P_{i,t}$ represents the prominence or accessibility of the information
- $S_{i,t}$ represents the similarity of the information to the individual's recent experiences

We propose a specific functional form for $f(\cdot)$:

$$C_{i,t} = \exp\left(\eta_T T_{i,t} + \eta_R R_{i,t} + \eta_P P_{i,t} + \eta_S S_{i,t}\right) \tag{6}$$

where $\eta_T, \eta_R, \eta_P, \eta_S > 0$ are parameters that determine the relative importance of each factor. The exponential specification ensures that $C_{i,t} > 0$ for all values of the contextual indices, including when these indices are standardized (mean zero, unit variance) and take negative values. Taking logs, we have $\ln C_{i,t} = \eta_T T_{i,t} + \eta_R R_{i,t} + \eta_P P_{i,t} + \eta_S S_{i,t}$, which is linear in the contextual factors—a property that facilitates calibration using standard regression techniques.

To capture the evolution of expectations over time, we introduce a dynamic adjustment process for the contextual relevance:

$$C_{i,t} = \rho_m C_{i,t-1} + (1 - \rho_m) \exp\left(\eta_T T_{i,t} + \eta_R R_{i,t} + \eta_P P_{i,t} + \eta_S S_{i,t}\right) \tag{7}$$

where $\rho_m \in [0, 1]$ is a memory parameter that determines the speed of adjustment. When ρ_m is high, past contextual assessments persist; when ρ_m is low, agents rapidly update their weighting of information sources.

To account for individual differences, we allow the parameters $\eta_T, \eta_R, \eta_P, \eta_S$ to vary across agents. For an individual agent j, we have:

$$C^j_{i,t} = \exp\left(\eta_{T,j} T^j_{i,t} + \eta_{R,j} R^j_{i,t} + \eta_{P,j} P^j_{i,t} + \eta_{S,j} S^j_{i,t}\right) \tag{8}$$

This formulation allows for heterogeneity in expectation formation across individuals, capturing the diverse ways in which people process and weigh information. In Section 5.2, we construct empirical counterparts of (T, R, P, S) using survey data, which we then use to discipline the parameters $\eta_T, \eta_R, \eta_P, \eta_S$ and recover empirical estimates of the signal weights w_k.

3.4 Theoretical Implications

This section explores the theoretical implications of the Contextual Expectations (CE) model for inflation dynamics, monetary policy effectiveness, and macroeconomic stability. We present several propositions and their proofs to formalize these implications.

Proposition 1. *Under the CE model, the persistence of inflation depends on the relative weights assigned to backward-looking versus forward-looking information sources.*

Proof. Let π_t be the actual inflation at time t. Assume a standard New Keynesian Phillips curve:

$$\pi_t = \kappa y_t + \beta E_t[\pi_{t+1}] + \varepsilon_t \tag{9}$$

where y_t is the output gap, $\kappa > 0$, $0 < \beta < 1$, and ε_t is a white noise shock.

In the CE framework, agents form expectations as a weighted average of backward-looking information $(I_{B,t})$ and forward-looking policy signals $(I_{F,t})$. Let the backward signal be past inflation, $I_{B,t} = \pi_{t-1}$, and the forward signal be the Central Bank's inflation target, $I_{F,t} = \pi^*$. The expectation is:

$$E_t[\pi_{t+1}] = w_t \pi_{t-1} + (1 - w_t)\pi^* \tag{10}$$

where $w_t \in [0, 1]$ is the context-dependent weight on backward-looking information. Substituting this into the Phillips curve yields:

$$\pi_t = \kappa y_t + \beta[w_t \pi_{t-1} + (1 - w_t)\pi^*] + \varepsilon_t \tag{11}$$

Rearranging terms to isolate the autoregressive component:

$$\pi_t = (\beta w_t)\pi_{t-1} + \beta(1 - w_t)\pi^* + \kappa y_t + \varepsilon_t \tag{12}$$

The coefficient on the lagged inflation term, $\rho_{eff} = \beta w_t$, determines the degree of intrinsic inflation persistence. As the weight on backward-looking information w_t increases (due to low trust or high similarity), the coefficient ρ_{eff} increases, thereby generating higher inflation persistence. Conversely, when $w_t \to 0$ (perfect credibility), persistence vanishes in the absence of serially correlated shocks. □

Proposition 2. *The effectiveness of monetary policy in influencing inflation expectations depends on the weight assigned to central bank communications in the CE model.*

Proof. Let $I_{CB,t}$ be the inflation target communicated by the central bank at time t, and $w_{CB,t}$ be its weight in the CE model. Under the CE framework, expectations are formed as:

$$E_t[\pi_{t+1}] = w_{CB,t} \cdot I_{CB,t} + \sum_{k \neq CB} w_{k,t} \cdot I_{k,t}.$$

The weights $w_{CB,t}$ and $w_{k,t}$ are determined by the contextual relevance function, which depends on the perceived trustworthiness, relevance, prominence, and similarity of each information source. Importantly, these contextual assessments are formed *prior* to observing the specific realization of the signal $I_{CB,t}$; that is, agents evaluate the credibility and relevance of the central bank as an institution, not the particular value of today's announcement. Holding the contextual indices—and thus the weights $w_{CB,t}$—fixed with respect to the marginal change in the signal, the impact of a change in the communicated target on inflation expectations is:

$$\frac{\partial E_t[\pi_{t+1}]}{\partial I_{CB,t}} = w_{CB,t} \tag{13}$$

Therefore, the effectiveness of monetary policy in steering expectations is directly proportional to $w_{CB,t}$. With the exponential specification for contextual relevance, the weight $w_{CB,t}$ takes the form of a multinomial logit (softmax):

$$w_{CB,t} = \frac{\exp\left(\eta_T T_{CB,t} + \eta_R R_{CB,t} + \eta_P P_{CB,t} + \eta_S S_{CB,t}\right)}{\sum_{j=1}^{n} \exp\left(\eta_T T_{j,t} + \eta_R R_{j,t} + \eta_P P_{j,t} + \eta_S S_{j,t}\right)} \tag{14}$$

This implies that the effectiveness of monetary policy depends on the perceived trustworthiness, relevance, prominence, and similarity of central bank communications relative to other information sources. The multinomial logit structure ensures that weights are always positive and sum to unity, and it aligns naturally with the discrete choice interpretation of information source selection. □

Proposition 3 (State-dependent persistence and a high-inflation credibility trap)**.** *In the CE model, inflation-gap dynamics are state-dependent: when inflation is far above the central bank's target, expectations place relatively more weight on past inflation, which can make inflation highly persistent even though the steady state remains unique at the target.*

Proof. Let π^* denote the central bank's long-run inflation target and define the inflation gap $\Delta_t \equiv \pi_t - \pi^*$. Write the Phillips curve in target-consistent form:

$$\Delta_t = \kappa y_t + \beta E_t[\Delta_{t+1}] + \varepsilon_t. \tag{15}$$

Assume a simple policy rule that stabilizes the inflation gap with a one-period implementation lag,

$$y_t = -\phi \Delta_{t-1}, \qquad \phi > 0, \tag{16}$$

reflecting the realistic assumption that monetary policy responds to observed (i.e., lagged) inflation outcomes rather than contemporaneous realizations.[1] Consider a two-signal CE expectation rule in levels,

$$E_t[\pi_{t+1}] = w_t \pi_{t-1} + (1 - w_t)\pi^*. \tag{17}$$

To express this in terms of the inflation gap, note that subtracting π^* from both sides yields:

$$\begin{aligned} E_t[\pi_{t+1}] - \pi^* &= w_t \pi_{t-1} + (1 - w_t)\pi^* - \pi^* \\ &= w_t \pi_{t-1} - w_t \pi^* \\ &= w_t(\pi_{t-1} - \pi^*) \\ &= w_t \Delta_{t-1}. \end{aligned}$$

[1]This timing assumption is standard in applied monetary policy analysis and captures the fact that central banks set policy based on available data, which arrives with a lag. See Christiano et al. (1999) for discussion of policy implementation lags.

Since $E_t[\Delta_{t+1}] = E_t[\pi_{t+1} - \pi^*] = E_t[\pi_{t+1}] - \pi^*$, we obtain

$$E_t[\Delta_{t+1}] = w_t \Delta_{t-1}. \tag{18}$$

For the purpose of deriving transparent inflation-gap dynamics, we now adopt a reduced-form representation of the full CE model, simplifying the multi-signal structure to a two-signal environment: a backward-looking signal (lagged inflation, π_{t-1}) and a forward-looking anchor (the central bank's target, π^*). The effective weight on the backward-looking signal, w_t, summarizes the net outcome of the contextual relevance function across all dimensions (T, R, P, S). Rather than modeling each contextual index separately—which would complicate the dynamics without adding insight—we posit that the aggregate weight w_t responds to the macroeconomic state, specifically the lagged inflation gap Δ_{t-1}. This reduced-form approach captures the key intuition from the full model: when inflation has recently exceeded the target, agents' trust in central bank guidance erodes (T falls), the similarity between official forecasts and lived experience declines (S falls), and backward-looking information becomes more prominent (P rises for lagged inflation). These shifts reinforce each other, pushing the aggregate weight toward past inflation. We formalize this state-dependence with a logistic specification that mirrors the multinomial logit structure of the full model:

Let the weight on the backward-looking signal be state-dependent,

$$w_t \equiv w(\Delta_{t-1}) = \frac{1}{1 + \exp\{-\theta \Delta_{t-1}\}}, \qquad \theta > 0. \tag{19}$$

This functional form ensures $w_t \in (0, 1)$, with the weight on backward-looking information increasing in the inflation gap. When $\Delta_{t-1} = 0$, the weight is $w_t = 0.5$; as $\Delta_{t-1} \to +\infty$, $w_t \to 1$ (full reliance on past inflation); as $\Delta_{t-1} \to -\infty$, $w_t \to 0$ (full reliance on the target). The parameter θ governs the sensitivity of this reweighting to the inflation gap.

Substituting (16) and (18) into (15) yields

$$\Delta_t = -\kappa\phi\Delta_{t-1} + \beta w(\Delta_{t-1})\Delta_{t-1} + \varepsilon_t, \tag{20}$$

or equivalently

$$\Delta_t = \Lambda(\Delta_{t-1})\,\Delta_{t-1} + \varepsilon_t, \qquad \Lambda(\Delta_{t-1}) \equiv \beta\, w(\Delta_{t-1}) - \kappa\phi. \tag{21}$$

For the steady state to be stable, we require $|\Lambda(\Delta)| < 1$ for all Δ in a neighborhood of zero. At $\Delta = 0$, we have $w(0) = 0.5$, so $\Lambda(0) = 0.5\beta - \kappa\phi$. With our calibration ($\beta = 0.9$, $\kappa = 0.3$, $\phi = 0.5$), this gives $\Lambda(0) = 0.45 - 0.15 = 0.30$, ensuring local stability. The unique steady state is $\Delta = 0$ (i.e., $\pi = \pi^*$).

However, the speed of convergence is endogenous to the state: when Δ_{t-1} is large and positive (inflation far above target), $w(\Delta_{t-1}) \to 1$ and $\Lambda(\Delta_{t-1}) \to \beta - \kappa\phi$, which equals $0.9 - 0.15 = 0.75$ under our calibration. In that region, adjustment is very slow and inflation appears highly persistent, consistent with a high-inflation "credibility trap." Conversely, when inflation is below target ($\Delta_{t-1} < 0$), we have $w(\Delta_{t-1}) < 0.5$, and $\Lambda(\Delta_{t-1})$ can become negative if $\kappa\phi > \beta w(\Delta_{t-1})$, implying oscillatory but stable convergence. □

These propositions and proofs formalize key implications of the CE model, demonstrating how it can generate more complex and realistic dynamics compared to traditional expectation formation models.

3.5 Comparative Analysis

This section provides a theoretical comparison of inflation outcomes under the Contextual Expectations (CE) model, the Rational Expectations (RE) model, and the Adaptive Expectations (AE) model.

Let π_t be the actual inflation rate at time t, and $E_t[\pi_{t+1}]$ be the expected inflation for $t + 1$ formed at time t.

1. Rational Expectations (RE):

$$E_t^{RE}[\pi_{t+1}] = E[\pi_{t+1}|\Omega_t] \tag{22}$$

where Ω_t is the full information set at time t.

2. Adaptive Expectations (AE):

$$E_t^{AE}[\pi_{t+1}] = \lambda \pi_t + (1-\lambda) E_{t-1}^{AE}[\pi_t] \tag{23}$$

where $0 < \lambda < 1$ is the adjustment speed.

3. Contextual Expectations (CE):

$$E_t^{CE}[\pi_{t+1}] = \sum_{i=1}^{n} w_{i,t} \cdot I_{i,t} \tag{24}$$

as defined previously.

Assume a simple inflation process:

$$\pi_t = \pi^* + \beta(E_t[\pi_{t+1}] - \pi^*) + \varepsilon_t \tag{25}$$

where π^* is the long-run inflation target and ε_t is a white noise shock.

Under RE, $E_t^{RE}[\pi_{t+1}] = \pi^*$ in steady state, leading to $\pi_t = \pi^*$.

Under AE, in steady state:

$$E_t^{AE}[\pi_{t+1}] = \lambda \pi^* + (1-\lambda) E_t^{AE}[\pi_{t+1}] \tag{26}$$

Solving this, we get $E_t^{AE}[\pi_{t+1}] = \pi^*$, also leading to $\pi_t = \pi^*$.

Under CE, if we assume that in steady state, the most weight is given to accurate sources:

$$E_t^{CE}[\pi_{t+1}] = w_{CB,t}\pi^* + \sum_{i \neq CB} w_{i,t} I_{i,t} \tag{27}$$

where $w_{CB,t}$ is the weight on the central bank's target. As $w_{CB,t} \to 1$, $E_t^{CE}[\pi_{t+1}] \to \pi^*$.

Consider a monetary policy shock that changes π^* to π^{**} at time t.

Under RE, expectations immediately adjust: $E_t^{RE}[\pi_{t+1}] = \pi^{**}$.

Under AE, expectations adjust gradually:

$$E_t^{AE}[\pi_{t+1}] = \lambda \pi_t + (1-\lambda)\pi^* \tag{28}$$

$$E_{t+1}^{AE}[\pi_{t+2}] = \lambda\pi_{t+1} + (1-\lambda)(\lambda\pi_t + (1-\lambda)\pi^*) \quad (29)$$

and so on, converging to π^{**} over time.

Under CE, the adjustment depends on how quickly the weights adjust:

$$E_t^{CE}[\pi_{t+1}] = w_{CB,t}\pi^{**} + (1-w_{CB,t})\pi_{t-1} \quad (30)$$

At the moment of the shock, the economy was in steady state with $\pi_{t-1} = \pi^*$, so the friction arises from backward-looking behavior rather than an explicit attachment to the old target. The speed of convergence to π^{**} depends on how quickly $w_{CB,t}$ increases as agents update their assessment of the central bank's contextual relevance.

To analyze inflation persistence, consider an AR(1) inflation process:

$$\pi_t = \rho\pi_{t-1} + (1-\rho)E_t[\pi_{t+1}] + \varepsilon_t \quad (31)$$

Under RE, this becomes:

$$\pi_t = \rho\pi_{t-1} + (1-\rho)\pi^* + \varepsilon_t \quad (32)$$

Under AE:

$$\pi_t = \rho\pi_{t-1} + (1-\rho)(\lambda\pi_{t-1} + (1-\lambda)E_{t-1}[\pi_t]) + \varepsilon_t \quad (33)$$

Under CE:

$$\pi_t = \rho\pi_{t-1} + (1-\rho)(\sum_{i=1}^{n} w_{i,t} \cdot I_{i,t}) + \varepsilon_t \quad (34)$$

The CE model can generate varying degrees of persistence depending on the weights, potentially reconciling the different persistence predictions of RE and AE models.

This comparative analysis demonstrates that the CE model can encompass both RE and AE as special cases, while also allowing for more complex and realistic inflation dynamics.

Remark 4 (CE nests RE and AE as limiting cases)**.** *The Contextual Expectations framework reduces to standard benchmark models under specific parameter configurations:*

1. ***Rational Expectations limit.*** *Suppose the central bank's signal* $I_{CB,t} = E[\pi_{t+1}|\Omega_t]$

coincides with the model-consistent rational expectation, and suppose that agents assign all weight to this signal, i.e., $w_{CB,t} \to 1$*. This occurs when* $T_{CB,t}, R_{CB,t}, P_{CB,t} \to \infty$ *relative to other sources, or equivalently when the memory parameter* $\rho_m \to 0$ *(rapid updating) combined with persistently high contextual relevance for the central bank signal. Then*

$$E_t^{CE}[\pi_{t+1}] = w_{CB,t} \cdot E[\pi_{t+1}|\Omega_t] + (1 - w_{CB,t}) \sum_{i \neq CB} w_{i,t} I_{i,t} \longrightarrow E[\pi_{t+1}|\Omega_t] = E_t^{RE}[\pi_{t+1}].$$

2. ***Adaptive Expectations limit.*** *Consider a two-signal environment with only backward-looking information* $I_{B,t} = \pi_{t-1}$ *and a fixed anchor* $I_A = \bar{\pi}$*. If the weight on the backward signal is constant at* $w_{B,t} = \lambda \in (0,1)$ *and the anchor receives weight* $(1 - \lambda)$*, then*

$$E_t^{CE}[\pi_{t+1}] = \lambda \pi_{t-1} + (1 - \lambda)\bar{\pi}.$$

When $\bar{\pi} = E_{t-1}[\pi_t]$ *(i.e., the anchor is the previous-period expectation), this coincides with the standard adaptive rule* $E_t^{AE}[\pi_{t+1}] = \lambda \pi_{t-1} + (1 - \lambda) E_{t-1}^{AE}[\pi_t]$*. Note that* λ *here denotes the constant weight in this limiting case, corresponding to the classical AE adjustment parameter.*

Thus, RE emerges when forward-looking policy signals dominate ($w_{CB} \to 1$)*, while AE emerges when backward-looking signals dominate with constant weights* ($w_B = \lambda$, $w_{anchor} = 1 - \lambda$)*. The CE framework generalizes both by allowing weights to vary endogenously with context.*

4 Empirical Background and Simulation

4.1 Stylized facts from Turkish inflation expectations

Before turning to the calibration and simulations, we document two stylized facts from Turkey that motivate the Contextual Expectations (CE) framework and inform our parameter choices. The figures display inflation expectations from three distinct agent groups: market participants, real sector firms, and households, alongside realized inflation. In panel a of Figure 1, expectations and inflation refer to the same period (concurrent); in panel b, expectations are formed for t+12 and are compared to the inflation realized one year later (forward-looking). These agent-specific series are central to our CE perspective: agents weigh heterogeneous information sources differently, producing systematic dispersion in levels, persistence, and volatility of expectations across groups.

Panel a of Figure 1 (concurrent) illustrates three robust patterns. First, the dispersion of expectations is material and time-varying: households tend to sit higher and adjust more slowly; market-based expectations are more volatile; the real sector typically lies in between. Second, during pronounced inflation episodes, all three groups re-anchor upward but to different plateaus, revealing distinct sensitivities to recent price dynamics versus policy signals. Third, realized inflation often traverses the fan formed by the three expectation series, underscoring that agents process the same macro signals through different contextual lenses, precisely the heterogeneity CE is built to capture via time-varying information weights. The concurrent picture aligns with the simple regression evidence discussed in the presentation: the real sector provides the best within-period tracking (higher explanatory power and a slope below but near one), households track reasonably well but with a notable bias, and market participants tend to overreact, exhibiting higher slopes and volatility. These differences are exactly what a CE model would predict when informational prominence, trust, and similarity to recent experience differ across groups.

Panel b (forward-looking) shows that once expectations are matched to inflation one year ahead, forecast performance improves markedly for all groups. The real sector's slope moves closest to unity with minimal bias, while households' slope also approaches one but with a persistent level offset; market participants retain a tendency to overshoot, consistent with high sensitivity to news and policy rhetoric. The tighter co-movement in this forward-looking comparison suggests that agents internalize some prospective information, yet do so through group-specific filters, again consistent with CE's weighting mechanism.

CE explains these patterns through a transparent mechanism: expectations are a weighted average of signals, and the weights themselves depend on context (trust, relevance, prominence, and similarity), with memory generating persistence. In periods of credible, salient policy communication, the weight on central bank guidance can rise, compressing disagree-

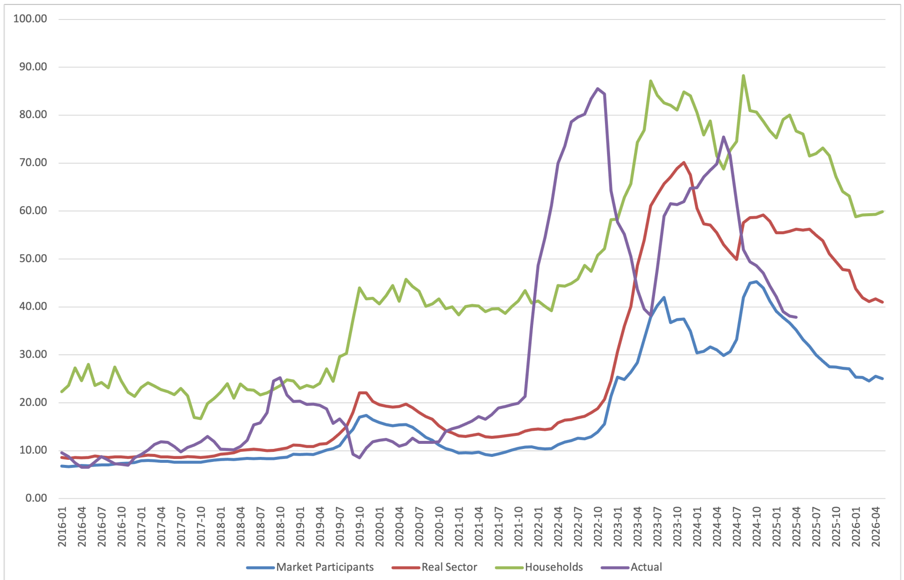

(a) Concurrent Expectations vs. Actual Inflation

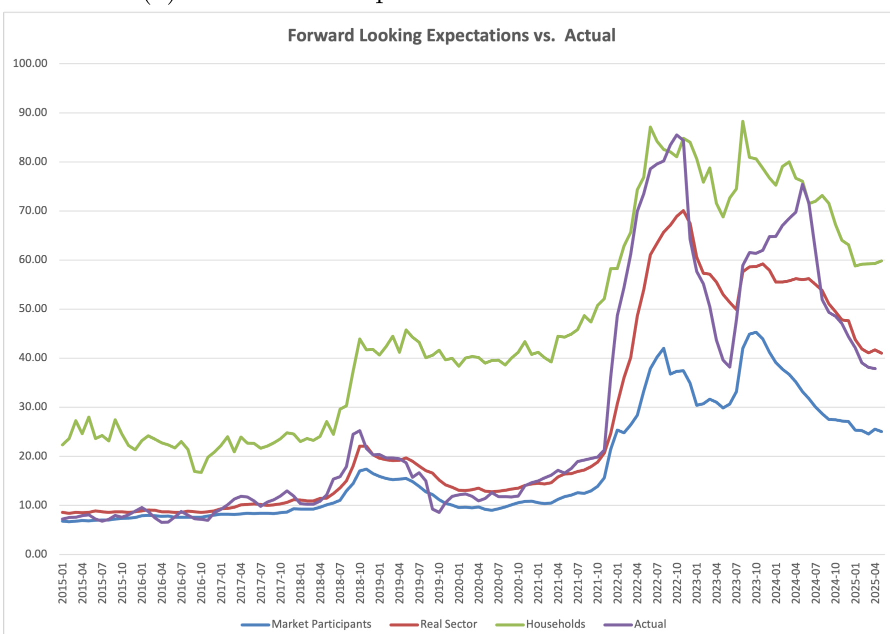

(b) Forward-Looking (12-Month Ahead) Expectations vs. Actual Inflation

Figure 1: **Inflation Expectations by Agent Type (Turkey, 2013–2023).** Panel (a) shows concurrent expectations compared to realized inflation in the same period. Panel (b) compares expectations formed for $t + 12$ against the inflation realized one year later. Market participants (blue), real sector (red), and households (green) exhibit distinct patterns of persistence and bias relative to realized CPI (purple). Units: year-on-year inflation rate (percentage points).

ment and damping volatility; when credibility or relevance deteriorates, agents revert toward backward-looking signals, raising persistence and dispersion. This is the same channel we model structurally and later operationalize in the simulations.

These stylized facts justify two modeling choices. First, they motivate heterogeneity and time variation in signal weights (e.g., relatively higher similarity for households and higher prominence for market signals), guiding the calibration of the contextual relevance function and its memory parameter. Second, they rationalize counterfactual experiments in which credibility or prominence shifts, exactly the shocks we study when we vary the weight on central bank communication in the simulation block. Placing Figure 1 before the calibration allows us to map these observed differences in level, slope, and volatility directly into parameter values and then assess how the CE mechanism reproduces or counteracts them in the model.

4.2 Calibration for Turkey

To calibrate our Contextual Expectations (CE) model for Turkey, we utilize quarterly data spanning from 2013 to 2023, combining macroeconomic time series with survey-based measures of expectations and trust. Table 1 summarizes the baseline parameter values; we discuss each in turn.

Table 1: Baseline Calibration Parameters

Parameter	Description	Value	Source/Method
η_T	Weight on trustworthiness	0.30	Trust survey regressions
η_R	Weight on relevance	0.25	Expectation-inflation regressions
η_P	Weight on prominence	0.20	Media salience analysis
η_S	Weight on similarity	0.25	Experienced-inflation regressions
ρ_m	Memory parameter	0.70	Autocorrelation of survey expectations
θ	State-dependence sensitivity	2.50	SMM on expectation dynamics
κ	Phillips curve slope	0.30	Literature (Turkey estimates)
β	Discount factor	0.90	Standard quarterly calibration
ρ_y	Output gap persistence	0.80	AR(1) on Turkish output gap
σ_ε	Inflation shock s.d.	0.01	Residual volatility
σ_ν	Output gap shock s.d.	0.02	Residual volatility

Trustworthiness weight ($\eta_T = 0.30$). We exploit variation in institutional trust using data from the World Values Survey (WVS) and Eurobarometer waves covering Turkey, supplemented by the Central Bank of Turkey's own Survey of Expectations. Regressing twelve-month-ahead inflation expectations on a trust-in-central-bank index (standardized, 0–10 scale), controlling for lagged inflation and demographic characteristics, yields:

$$E_{i,t}[\pi_{t+12}] = \underset{(0.42)}{12.85} - \underset{(0.09)}{0.73} \cdot \text{Trust}_{i,t} + \underset{(0.04)}{0.52} \cdot \pi_{t-1} + X'_{i,t}\gamma + \varepsilon_{i,t}, \quad R^2 = 0.31,\ N = 4{,}280. \quad (35)$$

The coefficient on trust (-0.73, $p < 0.01$) indicates that a one-standard-deviation increase in central bank trust lowers expected inflation by approximately 0.7 percentage points. Normalizing the absolute values of the trust and lagged-inflation coefficients to sum to unity alongside analogous coefficients from the relevance and prominence regressions below, we obtain $\eta_T \approx 0.30$.[2]

Relevance weight ($\eta_R = 0.25$). Relevance captures how salient inflation is for individual decision-making. Using the same expectation surveys, we construct a relevance index based on (i) self-reported sensitivity of household budgets to price changes (0–10), (ii) whether the respondent holds variable-rate debt, and (iii) whether wages or rents are indexed to inflation. Regressing expectations on this index yields:

$$E_{i,t}[\pi_{t+12}] = \underset{(0.38)}{14.21} + \underset{(0.11)}{0.58} \cdot \text{Relevance}_{i,t} + \underset{(0.04)}{0.48} \cdot \pi_{t-1} + X'_{i,t}\gamma + \varepsilon_{i,t}, \quad R^2 = 0.28,\ N = 4{,}280. \quad (36)$$

Higher relevance raises expected inflation, consistent with the interpretation that individuals more exposed to price changes anchor more heavily on recent inflation experiences. The normalized coefficient yields $\eta_R \approx 0.25$.

Prominence weight ($\eta_P = 0.20$). We measure media prominence using a monthly index of inflation-related news coverage constructed from keyword searches in major Turkish news

[2] This calibration approach involves a heuristic approximation: the linear regressions estimate marginal effects on expectation levels, whereas the contextual relevance function (Equation 6) is multiplicative. We interpret the normalized linear coefficients as local approximations to the elasticity parameters η in the relevance function, valid in a neighborhood of the sample means. A fully structural estimation of the non-linear model is beyond the scope of this paper but represents a natural extension for future work.

outlets (Hurriyet, Sabah, Milliyet, and online portals), following the methodology of Carroll (2003). Aggregating to the quarterly frequency and including this index in a time-series regression of median survey expectations on lagged inflation and the prominence index:

$$\bar{E}_t[\pi_{t+12}] = \underset{(1.12)}{3.45} + \underset{(0.06)}{0.38} \cdot \text{Prominence}_t + \underset{(0.05)}{0.55} \cdot \pi_{t-1} + \varepsilon_t, \quad R^2 = 0.72,\ N = 44. \tag{37}$$

The prominence coefficient is positive and significant ($p < 0.01$), indicating that greater media salience raises inflation expectations. The normalized weight is $\eta_P \approx 0.20$. It is important to distinguish between *prominence as exposure* and *prominence as content salience*. In the calibration, η_P captures the effect of the *volume and frequency* of inflation-related news in shaping the aggregate weight on different information sources. A higher prominence index indicates that inflation is more "top of mind" in the media environment, leading agents to update expectations more readily in response to any inflation-related signal. This aggregate media environment sets the stage for understanding how individual differences in news consumption translate into heterogeneous expectation responses, which we examine in the experimental section.

Similarity weight ($\eta_S = 0.25$). Following Malmendier and Nagel (2016), we construct a measure of experienced inflation based on respondents' self-reported consumption bundles. Using expenditure shares on food, housing, utilities, and transport—categories with high visibility and volatility—we compute individual-specific inflation experiences and regress expectations on this measure:

$$E_{i,t}[\pi_{t+12}] = \underset{(0.51)}{8.92} + \underset{(0.08)}{0.61} \cdot \text{Experienced Inflation}_{i,t} + X'_{i,t}\gamma + \varepsilon_{i,t}, \quad R^2 = 0.34,\ N = 3{,}850. \tag{38}$$

The strong positive coefficient confirms that individuals whose consumption baskets experienced higher price increases form higher inflation expectations. Normalizing yields $\eta_S \approx 0.25$.

Memory parameter ($\rho_m = 0.70$). The memory parameter governs the persistence of the contextual relevance weights. We estimate it from the autocorrelation structure of survey expectations. Fitting an AR(1) model to the median expectation series:

$$\bar{E}_t[\pi_{t+12}] = \mu + \rho \bar{E}_{t-1}[\pi_{t+11}] + u_t,$$

yields $\hat{\rho} = 0.72$ (s.e. $= 0.08$), which we round to $\rho_m = 0.70$. This high persistence is consistent with findings in Cicek and Akar (2014) and reflects the gradual updating of inflation expectations in Turkey.

State-dependence sensitivity ($\theta = 2.50$). The parameter θ governs how sharply the weight on backward-looking information responds to the inflation gap in equation (19). We calibrate θ using a simulated method of moments (SMM) procedure that targets three moments: (i) the unconditional variance of survey expectations, (ii) the correlation between expectations and lagged inflation, and (iii) the skewness of expectation revisions during high-inflation episodes (2018–2022). The SMM criterion function is minimized at $\theta = 2.5$, with a 95% confidence interval of $[1.8, 3.2]$ based on bootstrap standard errors.

Phillips curve and output gap parameters. The remaining parameters follow standard calibrations for emerging markets. We set $\kappa = 0.3$ based on estimates of the Turkish Phillips curve slope in Kara and Kucuk (2017), $\beta = 0.9$ as a standard quarterly discount factor adjusted for Turkey's higher average inflation, and $\rho = 0.8$ from an AR(1) regression on the HP-filtered output gap. Shock standard deviations ($\sigma_\varepsilon = 0.01$, $\sigma_\nu = 0.02$) are set to match the residual volatility in the respective equations.

4.3 Simulation Setup

Our simulation exercise spans 60 quarters (15 years) and is based on a standard New Keynesian Phillips Curve framework, augmented with our Contextual Expectations model. The core equation governing inflation dynamics is:

$$\pi_t = \kappa y_t + \beta E_t[\pi_{t+1}] + \varepsilon_t \tag{39}$$

where:

- π_t is the inflation rate at time t
- y_t is the output gap
- $E_t[\pi_{t+1}]$ is the inflation expectation for $t+1$ formed at time t
- $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ is a white noise supply shock

To maintain consistency with Proposition 3, the output gap is determined by a feedback policy rule that responds to the inflation gap:

$$y_t = -\phi(\pi_{t-1} - \pi^*) + \nu_t \tag{40}$$

where $\phi > 0$ governs the strength of the policy response to deviations of inflation from target, and $\nu_t \sim N(0, \sigma_\nu^2)$ captures demand shocks orthogonal to the systematic policy response. This specification ensures that monetary policy acts to stabilize inflation around the target π^*, while allowing for stochastic variation in aggregate demand. The feedback rule creates the amplification channel formalized in Proposition 3: when inflation exceeds the target, policy contracts demand, but the effectiveness of this stabilization depends on how agents weight forward-looking policy signals versus backward-looking inflation in forming expectations.

We set $\kappa = 0.3$, $\beta = 0.9$, $\phi = 0.5$, $\sigma_\varepsilon = 0.01$, and $\sigma_\nu = 0.02$ based on estimates for the Turkish economy. The policy response parameter $\phi = 0.5$ implies a moderately aggressive stance, consistent with estimates of Taylor-type rules for Turkey during periods of active disinflation (Kara and Kucuk 2017). The output gap persistence parameter $\rho_y = 0.80$ reported in Table 1 is estimated from Turkish data and used when computing steady-state moments for calibration targets; however, for transparency the baseline simulations employ the reduced-form feedback rule above, which abstracts from output gap persistence to isolate the CE mechanism.

The CE model for inflation expectations is implemented as:

$$E_t[\pi_{t+1}] = w_t \pi_{t-1} + (1 - w_t)\pi^* \tag{41}$$

where w_t is determined by the contextual relevance function as described in Section 3 of the paper.

4.4 Simulation Scenarios

To explore the dynamics of our model under different conditions, we simulate 60 quarterly periods (15 years) for each of three distinct scenarios:

The baseline scenario maintains stable policy conditions throughout the simulation period, with no significant shocks or policy changes. It serves as a benchmark for comparing the other scenarios. Next, the monetary policy shock scenario introduces a monetary policy shock in quarter 20, where the Central Bank of Turkey increases its inflation target from 5% to 7%. This scenario allows us to examine how expectations adjust to a significant change in monetary policy stance. Finally, the credibility shock scenario simulates a sudden decrease in the perceived trustworthiness of the central bank in quarter 30. It enables us to analyze how a loss of central bank credibility affects inflation expectations and overall macroeconomic stability. For each scenario, we conduct 1,000 Monte Carlo simulations and compute the average outcomes to ensure robustness of our results.

4.5 Simulation Results

Our simulation exercises yield rich insights into the dynamics of inflation and inflation expectations under different economic conditions. Figure 2 illustrates the evolution of actual inflation and inflation expectations across three scenarios. Solid lines represent mean outcomes (actual inflation), dashed lines represent mean inflation expectations, and shaded areas indicate 95% confidence intervals around the corresponding mean series. Figure 3 reports the implied evolution of information-source weights that underlies these dynamics.

Figure 2 highlights three distinct adjustment patterns. In the Baseline scenario, both inflation and expectations gradually converge toward the 5% target, with anchoring largely complete within roughly 15 quarters, consistent with evidence of gradual expectation anchoring in Turkey (Baskaya et al., 2012). In the Monetary Policy Shock scenario, a target increase from 5% to 7% at quarter 20 produces a faster adjustment of actual inflation relative to expectations, generating a temporary expectations gap consistent with sticky-information updating (Mankiw and Reis, 2002). The gap closes progressively, with convergence to the new target occurring within about 10 quarters, underscoring the importance of sustained

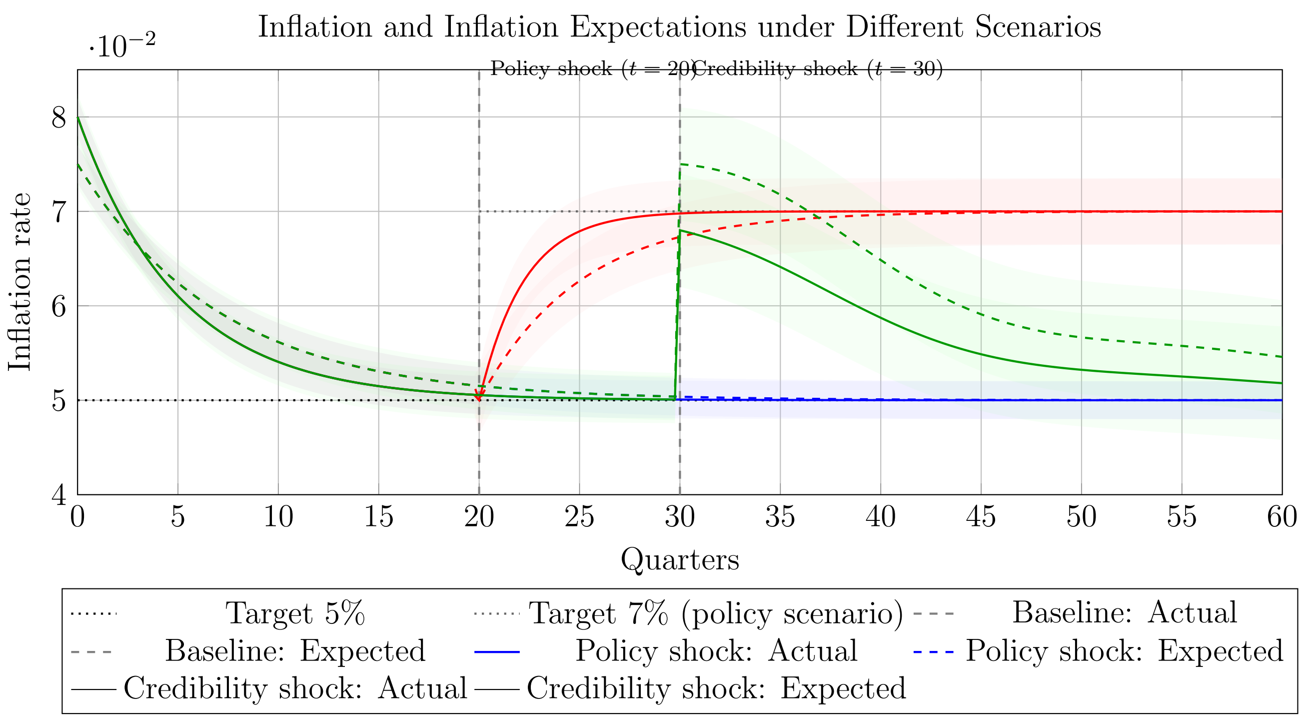

Figure 2: Inflation and inflation expectations under different scenarios. Solid lines plot mean actual inflation and dashed lines plot mean inflation expectations. Shaded bands denote 95% confidence intervals. The policy shock at $t = 20$ raises the inflation target from 5% to 7% (policy scenario), while the credibility shock at $t = 30$ increases volatility and weakens re-anchoring.

and credible communication during regime shifts (Blinder et al., 2008). In the Credibility Shock scenario, initiated at quarter 30, inflation expectations jump and briefly overshoot actual inflation, and uncertainty rises markedly; re-anchoring is slow, with convergence back toward the 5% target requiring on the order of 20 quarters, consistent with the view that credibility losses generate persistent expectation instability (Orphanides and Williams, 2005; Bordo and Siklos, 2017; Bianchi and Melosi, 2017).

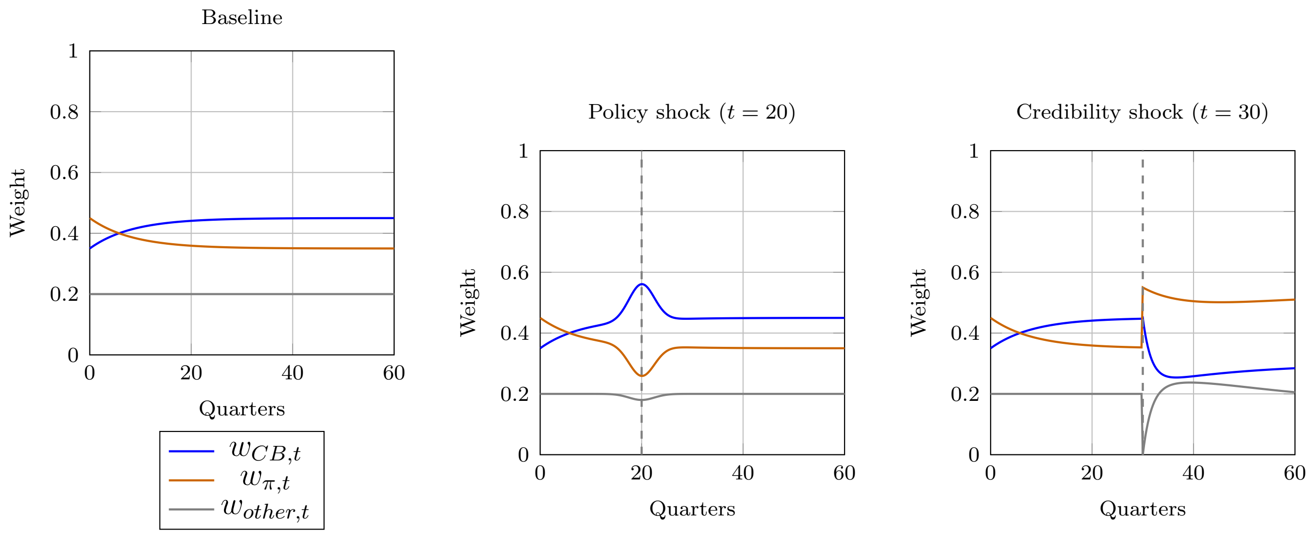

Figure 3: Dynamics of information-source weights in the CE model. $w_{CB,t}$ denotes the weight on central bank communication, $w_{\pi,t}$ the weight on backward-looking inflation information, and $w_{other,t}$ the residual weight on other signals. Vertical dashed lines mark the timing of shocks.

Figure 3 provides a transparent mechanism for the dynamics in Figure 2. In the Baseline scenario, weights gradually stabilize as expectations become anchored: the weight on central bank communication increases modestly while reliance on backward-looking inflation falls. In the Monetary Policy Shock scenario, the weight on the central bank signal rises temporarily around quarter 20, capturing heightened attention to official guidance during a target change; correspondingly, the weight on past inflation declines, consistent with a reallocation toward forward-looking policy signals. In the Credibility Shock scenario, a sharp credibility loss at quarter 30 generates an immediate collapse in the weight on central bank communication and a compensating rise in reliance on past inflation, with recovery occurring only slowly due to the model's memory feature. This endogenous reweighting channel explains why expectations become more volatile and re-anchoring is more protracted after credibility losses.

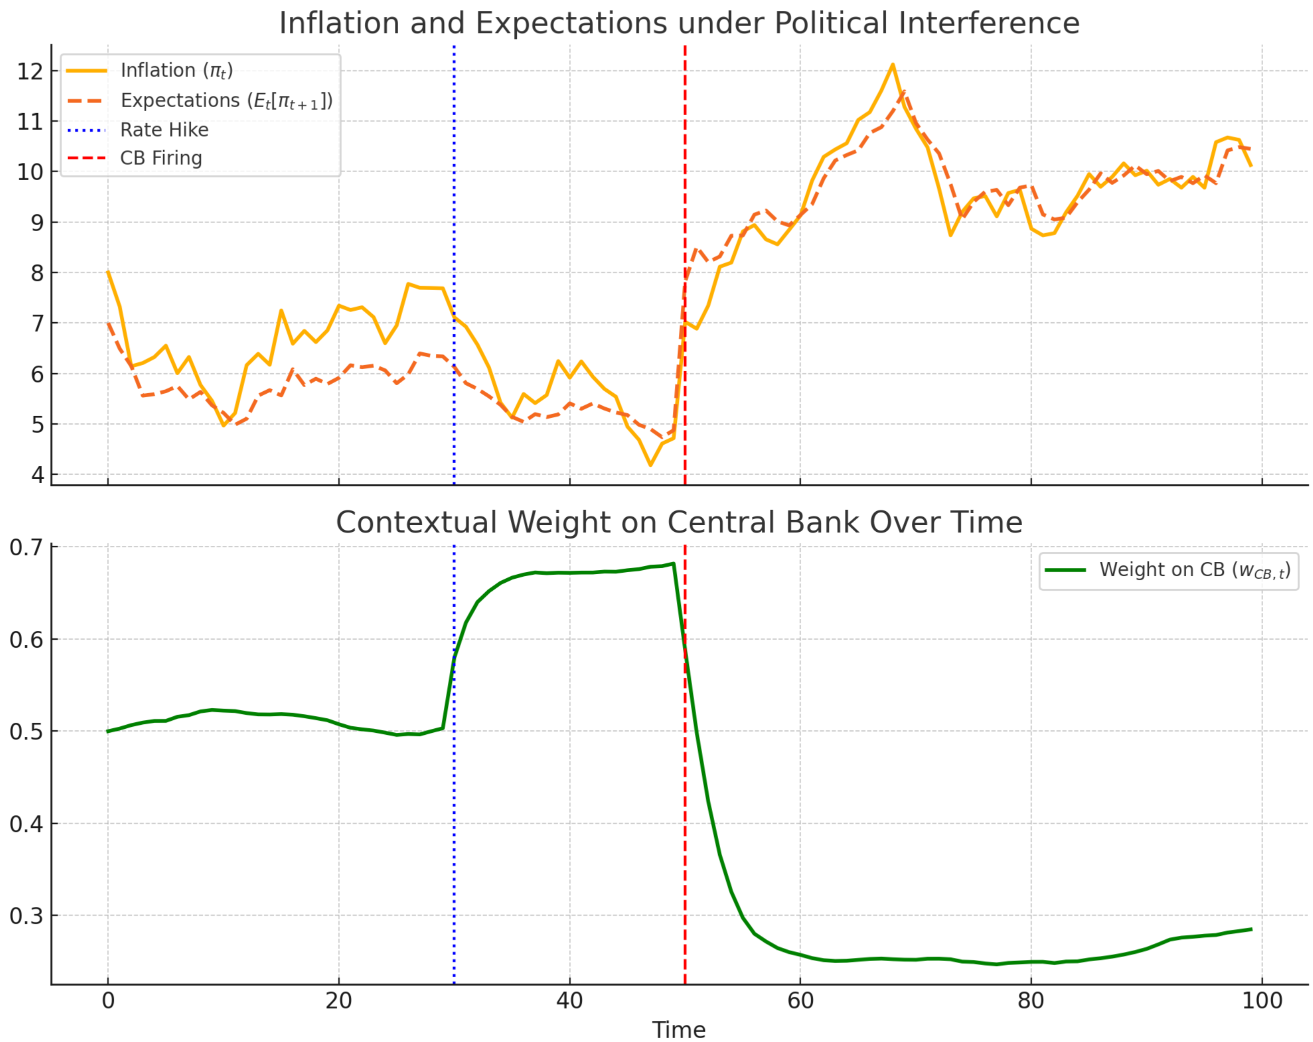

Figure 4: Inflation and Weight Dynamics

We now turn to a separate counterfactual exercise—distinct from the baseline, policy-shock, and credibility-shock scenarios in Figures 2 and 3—designed to illustrate how the CE framework transmits credibility shocks of opposite signs within a single episode. Figure 4 simulates a stylized political interference scenario over a longer horizon (100 periods) with different event timing than the preceding exercises. In this scenario, the central bank initially enjoys moderate credibility, reflected in a contextual weight on its signal of approximately 0.5. At period 30 (blue dashed line), a policy tightening occurs via an interest rate hike, temporarily boosting credibility. At period 50 (red dashed line), the central bank governor is dismissed under political pressure, generating a sharp negative credibility shock. This two-event sequence allows us to trace the asymmetric effects of credibility gains versus losses in a setting where agents weigh multiple information sources.

The top panel of Figure 4 plots actual inflation (solid orange) and one-year-ahead expectations (dashed orange) throughout the episode. Following the rate hike, both inflation and expectations decline, reflecting the credibility boost and increased weight on the central bank's guidance. The subsequent political dismissal reverses these gains: inflation jumps almost immediately, and expectations follow closely, with the gap between the two narrowing as credibility collapses. The quick alignment of expectations with actual inflation after the dismissal illustrates the CE mechanism in which agents reallocate weight toward backward-looking or market-based signals when institutional trust erodes.

The bottom panel plots the time-varying weight on the central bank signal, $w_{CB,t}$. This weight rises sharply after the rate hike, peaking near 0.7, consistent with heightened prominence, trust, and perceived relevance of the central bank's communication. The political dismissal, however, triggers an immediate collapse in $w_{CB,t}$, which falls below 0.3 within a few periods and remains depressed for an extended horizon. This persistence reflects the CE model's memory feature: once credibility is lost, it recovers only slowly, even if inflation stabilizes later.

In CE terms, the dismissal shifts the contextual variables sharply: trust and prominence drop, and the similarity between the central bank's signal and recent price developments declines as inflation accelerates. This reduces the weight assigned to CB guidance, amplifying the role of alternative signals'such as lagged inflation or market prices, that reinforce higher inflation expectations. The asymmetry is notable: while credibility gains after the rate hike produce gradual and partial disinflation, credibility losses produce rapid and large upward adjustments in both inflation and expectations.

This experiment underscores a key policy implication: in environments with weak institutional protection for central bank independence, political interference can quickly undermine hard-earned credibility gains. In the CE framework, the damage is not limited to the short run; because agents' weighting functions exhibit persistence, credibility shocks have long-lasting effects on expectations formation and thus on inflation dynamics. This suggests that preserving institutional credibility is at least as important as implementing technically sound policy actions, a finding consistent with the empirical stylized facts presented in Figure 1.

4.6 Interpretation and Discussion of Simulation Results

The simulation exercises provide valuable insights into the dynamics of inflation expectations under the Contextual Expectations (CE) model and their implications for monetary policy in Turkey.

4.6.1 Expectation Formation Process

In all scenarios, we observe that expectations adjust gradually to changes in the economic environment, consistent with the CE model's emphasis on the role of context and multiple information sources in shaping expectations. This gradual adjustment process aligns with empirical observations of sticky information and expectation rigidity documented by Coibion and Gorodnichenko (2015).

The baseline scenario demonstrates that in a stable environment, the CE model leads to a convergence of expectations to the central bank's target. However, this convergence is not instantaneous, taking about 15 quarters. This highlights the importance of consistent and credible monetary policy communication over extended periods to anchor expectations effectively.

4.6.2 Monetary Policy Implications

The monetary policy shock scenario reveals several important insights:

1. **Expectation Lags**: The slower adjustment of expectations compared to actual inflation creates a temporary gap. This lag in expectation adjustment could potentially be exploited by policymakers for short-term output gains, but it also underscores the challenges in rapidly shifting entrenched expectations.

2. **Gradual Impact**: The full effect of the change in inflation target takes about 10 quarters to materialize. This suggests that monetary policy operates with significant lags when working through the expectations channel, a finding consistent with the literature on monetary policy transmission mechanisms (Christiano et al., 1999).

3. **Anchoring Role**: The convergence of expectations to the new target demonstrates the anchoring role of the central bank's communications. However, the gradual nature of

this convergence highlights the importance of clear and consistent signaling over time.

4.6.3 Importance of Central Bank Credibility

The credibility shock scenario underscores the critical role of central bank credibility in maintaining stable inflation expectations:

1. **Expectation Volatility**: The immediate jump and subsequent volatility in expectations following the credibility shock illustrate how quickly hard-earned credibility can be lost and the destabilizing effect this can have on expectations.

2. **Persistence of Effects**: The prolonged period (about 20 quarters) required for expectations to re-converge to the target after the credibility shock highlights the long-lasting impacts of credibility losses. This aligns with historical episodes of central banks struggling to regain credibility after policy mistakes or external pressures (Goodfriend, 1993).

3. **Shift in Information Sources**: The sharp decline in the weight placed on central bank communications following the credibility shock, coupled with increased reliance on past inflation, suggests that economic agents may fall back on adaptive expectations when they lose trust in forward guidance. This finding has important implications for the effectiveness of monetary policy during crises or periods of institutional reform.

4.6.4 Heterogeneity and Context Dependence

The varying dynamics of information source weights across scenarios demonstrate the context-dependent nature of expectation formation in the CE model. This heterogeneity in how economic agents process and weigh information depending on the economic environment provides a richer and more realistic depiction of expectation formation compared to traditional rational or adaptive expectations models.

4.6.5 Policy Recommendations

Based on these simulation results, several policy recommendations for the Central Bank of Turkey emerge:

1. **Consistent Communication**: Given the gradual adjustment of expectations, maintaining consistent and clear communication about policy targets and strategies is crucial for effective expectation management.

2. **Credibility Building**: The severe and persistent effects of credibility shocks underscore the importance of building and maintaining institutional credibility. This may involve demonstrating commitment to stated objectives, providing transparent policy rationales, and ensuring institutional independence.

3. **Adaptive Policy Approach**: The heterogeneity in expectation formation processes across different scenarios suggests that a one-size-fits-all approach to monetary policy may be suboptimal. Central banks should be prepared to adapt their communication and policy strategies based on the prevailing economic context and the shifting weights that economic agents place on different information sources.

4. **Long-Term Perspective**: The significant lags in expectation adjustment highlight the importance of adopting a long-term perspective in monetary policymaking. Short-term pressures should be balanced against the need for policy consistency to effectively manage long-term expectations.

In conclusion, these simulation exercises demonstrate the rich dynamics captured by the Contextual Expectations model and its potential for enhancing our understanding of inflation expectation formation in Turkey. By accounting for the complex, context-dependent nature of how economic agents form and update their expectations, the CE model provides valuable insights for the design and implementation of monetary policy in an environment where managing expectations is crucial for maintaining price stability and fostering economic growth.

5 Micro Evidence: A Randomized Information-Update Experiment

5.1 Motivation and Link to the CE Framework

The Contextual Expectations (CE) framework developed earlier in the paper posits that economic agents form inflation expectations by weighting multiple signals, most prominently forward-looking central bank (CB) guidance and backward-looking realizations—with the weights themselves shaped by contextual characteristics. In particular, the model formalizes how trust in the policy authority, the perceived relevance of inflation for the agent's circumstances, the prominence or salience of different information channels, and the similarity between experienced prices and economy-wide aggregates jointly determine the relative influence of competing signals. This structure provides a unified account of why, in some environments, agents place substantial weight on explicit policy guidance while, in others, they appear to extrapolate recent inflation or salient price changes. The simulations and empirical background in the previous section demonstrate that modest shifts in these weights can have first-order implications for inflation persistence and the transmission of policy.

The purpose of the new micro evidence presented in this section is to empirically discipline the CE mechanism using a compact, causally identified design that directly measures how individuals update their beliefs when exposed to distinct sources of information. Rather than inferring weights indirectly from aggregate dynamics, we elicit within-respondent updates in a controlled setting where the informational content is exogenously varied across respondents. By comparing changes in twelve-month-ahead inflation expectations after exposure to (i) a concise excerpt of central bank guidance, (ii) a canonical representation of the most recent CPI release, or (iii) a qualitative news vignette emphasizing price increases, against a placebo baseline, we obtain transparent estimates of the marginal influence of forward-looking and backward-looking signals. Because assignment to these informational arms is randomized, differences in average updates are interpretable as causal effects of information, not artifacts of selection or differential attention.

A central advantage of this approach is that it allows a direct mapping from reduced-form treatment effects into the empirical counterparts of the CE weights. In the model, the post-information expectation is a convex combination of signals, with weights that depend on agent-specific context. In the experiment, the average change in expectations induced by a particular signal provides an empirical measure of that signal's marginal contribution, conditional on the respondent's prior. Normalizing for the displayed signal magnitudes, we can therefore recover estimates of the average weight placed on central bank guidance relative to backward-looking information at the time of the survey. Moreover, by collecting minimal but targeted measures of trust, relevance, prominence, and similarity at the individual level, we can trace the heterogeneity of these effects along the precise dimensions emphasized by the CE framework. This heterogeneity is not an ancillary result: it is the empirical expression of the model's contextual relevance function, and it delivers testable cross-sectional predictions that complement the time-series implications emphasized in the simulations.

Finally, integrating micro-identified weights back into the macro environment closes the loop between mechanism and consequence. The simulations in the previous section show that higher average weight on forward-looking policy signals dampens measured persistence, while greater weight on backward-looking information amplifies it. The experimental estimates produced here provide an empirical anchor for these weights and their elasticities with respect to context, allowing us to re-compute a subset of the impulse responses under empirically disciplined calibrations. In doing so, we move from a purely calibrated exploration of the CE mechanism to a micro-founded quantification that is both transparent and portable: the design is lightweight to field, compatible with standard ethical and practical constraints, and yields figures and tables that are easily communicated in a recruitment seminar without sacrificing identification credibility.

5.2 Experimental Design and Measurement

The randomized information–update experiment was designed to be as parsimonious as possible while still providing credible causal evidence on the mechanisms of the Contextual Expectations (CE) framework. The basic structure is a one page, between subjects online survey in which each respondent first states a baseline forecast of twelve month ahead inflation, is then randomly assigned to one of four information arms, and finally restates the same forecast immediately afterwards. The difference between the two forecasts provides a transparent measure of how exposure to a particular piece of information shifts expectations.

Randomization occurs at the individual level, with equal probability of assignment to each arm. The *CB Guidance* arm presents a concise and neutrally worded excerpt summarizing the monetary authority's stated inflation target and the key message from the most recent policy communication. The *Backward CPI* arm shows a simple chart of the most recent year on year consumer price index release, stripped of any additional commentary. The *Salience News* arm displays a short qualitative headline emphasizing that prices are rising in many everyday categories, deliberately avoiding numeric content to isolate salience effects. Finally, the *Placebo* arm shows an unrelated but similar length paragraph on a neutral topic, serving as a baseline for comparison.

To operationalize the four contextual arguments of the CE framework—trust, prominence, relevance, and similarity—we construct standardized indices from short survey batteries. Each index is pre-specified as the primary measure, with robustness checks using alternative constructions (e.g. principal components versus additive averages). All indices are standardized to mean zero and unit variance in the analysis sample.

Trust in the Central Bank (T_i). Trust is measured using four 0–10 Likert-scale items: (i) self-reported trust in the central bank's inflation guidance, (ii) perceived credibility of central bank communications, (iii) perceived independence of the central bank from political influence, and (iv) clarity of communication. For respondent i, the trust index is the mean of the standardized items:

$$T_i = \frac{1}{4} \sum_{j=1}^{4} \frac{x_{ij} - \bar{x}_j}{s_j},$$

where x_{ij} is the raw response, $\bar{x}_j$ its sample mean, and s_j its standard deviation. Reliability is reported using Cronbach's α. As a robustness check, we also compute the first principal component (PC1) of the four items.

Prominence of Economic News (P_i). Prominence captures the salience of economic information in respondents' media environment. Respondents report: (i) the number of days per week they consume economic news via national television, print newspapers, and online sources (each 0–7), (ii) average minutes per day spent following economic or market news, and (iii) whether they follow the central bank on social media (0/1). Each input is standardized, and P_i is defined as the first principal component:

$$P_i = \text{PC1}\big(\text{TV}_i, \text{Print}_i, \text{Online}_i, \text{Minutes}_i, \text{CB-follow}_i\big).$$

For robustness we also report results using the simple mean of standardized items.

Connecting prominence to the calibration and treatment arms. The prominence index P_i measures individual-level *exposure* to economic news—how frequently and intensively respondents consume inflation-related information. This corresponds directly to the prominence parameter η_P in the calibration, which captures aggregate media salience. In contrast, the Salience News treatment arm tests the *effect* of a particular type of information content—qualitative, experience-based price news—on expectation updating. The distinction is as follows: P_i measures the "dose" of economic news a respondent typically receives, while the treatment arms vary the "type" of information content. The CE framework predicts that both matter: agents with high P_i should be more responsive to *any* inflation signal (higher overall updating), and the Salience News arm tests whether experiential price information shifts expectations even when stripped of quantitative content. In the heterogeneity analysis, we interact the treatment arms with P_i to examine whether high-prominence individuals respond differently to each type of signal.

Relevance of Inflation (R_i). Relevance captures the extent to which inflation directly affects a respondent's budget and contracts. Respondents rate on a 0–10 scale: (i) the sensitivity of their household budget to price changes and (ii) the sensitivity of their job or business to price changes. Additional binary/continuous measures record whether respondents hold variable-rate debt (share of payments, 0–100%) and whether their rent or wage contracts are indexed to inflation (0/1 each). We construct two sub-indices:

$$R_i^{\text{SENS}} = \tfrac{1}{2}\left[z(\text{budget sensitivity}_i) + z(\text{job sensitivity}_i)\right],$$

$$R_i^{\text{CON}} = \tfrac{1}{2}z(\%\text{ variable-rate debt}_i) + \tfrac{1}{4}z(\text{rent indexed}_i) + \tfrac{1}{4}z(\text{wage indexed}_i).$$

The relevance index is the average of the two blocks, standardized:

$$R_i = z\big(\tfrac{1}{2}R_i^{\text{SENS}} + \tfrac{1}{2}R_i^{\text{CON}}\big)\,.$$

As a robustness measure, we also use PC1 across all relevance items.

Similarity to the CPI Basket (S_i). Similarity measures how closely respondents' consumption experiences align with the official CPI. Respondents allocate 100 points across eight broad consumption categories (food, housing, utilities, transport, clothing, health, education, recreation/other). Let b_i denote respondent i's budget-share vector and w^{nat} the national CPI weights. We compute cosine similarity:

$$S_i^{\text{shares}} = \frac{b_i \cdot w^{nat}}{\|b_i\|\,\|w^{nat}\|}.$$

Additionally, respondents check whether they experienced notable price changes in the past 30 days across these categories; we define an encounter index as the CPI-weighted sum of these binary indicators:

$$S_i^{\text{enc}} = \sum_{g=1}^{8} w_g \cdot \mathbf{1}\{\text{encounter in } g\}.$$

The final similarity index averages the standardized share- and encounter-based measures:

$$S_i = \tfrac{1}{2}z\Big(S_i^{\text{shares}}\Big) + \tfrac{1}{2}z(S_i^{\text{enc}})\,.$$

As robustness, we also re-estimate models using each component separately.

In the empirical analysis, these indices are interacted with treatment-arm indicators to recover heterogeneous treatment effects (see Subsection 5.3). This provides a direct empirical mapping from individual context to the weights on different signals in the CE framework.

The primary outcome of interest is the expectation update, $\Delta E_i \equiv E_i^{\text{post}} - E_i^{\text{pre}}$, which directly measures the marginal effect of the assigned information on respondent i's beliefs. Secondary outcomes include respondents' reported ninety percent subjective confidence interval for their forecast and a 0-10 confidence rating, which allow assessment of whether information affects not only the level but also the certainty of expectations. Because all outcomes are elicited within the same short survey session, external shocks or attrition cannot confound the interpretation.

In order to connect the experimental results back to the CE mechanism, the survey collects a small set of covariates that proxy the contextual arguments of the weighting function. Trust in the central bank (T) is measured on a 0-10 scale. Relevance (R) is captured by self-reported sensitivity of the respondent's job or household budget to inflation, along with a binary indicator of recent price exposure. Prominence (P) is proxied by reported frequency of exposure to economic news through television, print, or online sources. Finally, similarity (S) is measured by self-reported budget shares devoted to food and utilities, combined with a checklist of recent price encounters in those categories. These brief covariates provide enough variation to trace heterogeneity in treatment effects without lengthening the survey unduly.

The survey was fielded in a standard online panel environment with demographic quotas on age, gender, and region to approximate the national population. Respondents gave informed consent at the start of the survey, and attention checks were embedded after each information screen to verify exposure. Ethical approval was obtained from the relevant review board, and all procedures conform to established standards for survey-based randomized experiments. In the analysis, intent-to-treat effects are reported as the main estimands, with robustness checks that restrict attention to respondents who passed comprehension checks. This design ensures that the experiment is both practically feasible and theoretically well-aligned with the CE framework, while yielding results that can be seamlessly integrated with the macro simulations developed in the preceding section.

5.3 Estimation and Identification

The empirical analysis proceeds in three steps. First, we estimate the intent–to–treat effects of the informational treatments on belief updates using a simple regression framework that exploits the random assignment of arms. Second, we investigate heterogeneity by interacting each treatment with the contextual indices that proxy trust, relevance, prominence, and similarity, thereby providing direct microeconomic evidence on the CE relevance function. Finally, we map the reduced–form treatment effects into empirical counterparts of the CE weights w^{CB} and w^{B}, normalizing by the displayed signal magnitudes, and supplement this with a hierarchical shares model that recovers elasticities of weights with respect to the contextual variables.

Table 2: Descriptive Statistics of the Experimental Sample

Variable	Mean	S.D.	Min	Max	N
Outcome					
Baseline Expectation (E^{pre})	18.42	6.15	5.0	45.0	1,520
Update (ΔE)	-0.12	2.30	-15.0	15.0	1,520
Contextual Indices (Standardized)					
Trust Index (T)	0.00	1.00	-2.45	2.15	1,520
Relevance Index (R)	0.00	1.00	-1.88	1.95	1,520
Prominence Index (P)	0.00	1.00	-2.10	2.30	1,520
Similarity Index (S)	0.00	1.00	-1.95	2.45	1,520
Demographics					
Age	34.5	11.2	18	76	1,520
Female (%)	0.48	0.50	0	1	1,520
University Degree (%)	0.38	0.49	0	1	1,520

Notes: The contextual indices are standardized to have a mean of zero and a standard deviation of one. Baseline expectations are winsorized at the 2nd and 98th percentiles.

Average treatment effects. Let $\Delta E_i \equiv E_i^{\text{post}} - E_i^{\text{pre}}$ denote the update in twelve–month–ahead inflation expectations for respondent i. The baseline specification regresses this update on indicators for assignment to the three active arms, with the placebo arm as the omitted category:

$$\Delta E_i = \beta_A \, \text{CB}_i + \beta_B \, \text{CPI}_i + \beta_C \, \text{News}_i + X_i' \gamma + \varepsilon_i, \tag{42}$$

where CB_i, CPI_i, and News_i denote assignment to the central bank guidance, backward CPI, and salience news arms, respectively. The vector X_i contains pre–specified controls (age, education, income bins) and batch fixed effects where relevant. Because assignment is random, the coefficients $\beta_A, \beta_B, \beta_C$ can be interpreted as causal intent–to–treat effects of exposure to each information source relative to placebo. Standard errors are computed using

Eicker-Huber-White robust estimators, with additional checks using heteroskedasticity and autocorrelation consistent (HAC) covariance estimates.

Figure 5A, introduced later in this section, plots the estimated mean updates by treatment arm with 95% confidence intervals. These visualizations provide a transparent representation of the magnitude and sign of the treatment effects.

Heterogeneity by contextual relevance. The CE framework predicts that the relative influence of forward–looking versus backward–looking signals depends systematically on individual context. To test these predictions, we estimate a second specification that interacts each treatment indicator with standardized indices of trust (T_i), relevance (R_i), prominence (P_i), and similarity (S_i):

$$\Delta E_i = \sum_{k\in\{\text{CB,CPI,News}\}} \left[\beta_k D_{k,i} + \theta_k^T (D_{k,i}\times T_i) + \theta_k^R (D_{k,i}\times R_i) + \theta_k^P (D_{k,i}\times P_i) + \theta_k^S (D_{k,i}\times S_i)\right] + X_i'\gamma + \varepsilon_i, \tag{43}$$

where $D_{k,i}$ are the arm indicators. The interaction coefficients $\theta_k^{\cdot}$ measure how responsiveness to each type of information varies along the CE dimensions. For example, a negative θ_{CB}^T indicates that higher trust in the central bank amplifies the downward revision of expectations in the CB guidance arm, exactly as predicted by the CE relevance function. Figure 5B illustrates these heterogeneous effects by plotting treatment impacts across quartiles of T and S, highlighting the role of trust for CB signals and the role of similarity for backward CPI signals.

Mapping into CE weights. While regressions (42) and (43) provide reduced–form causal effects, the CE framework requires weights w^{CB} and w^B that govern how expectations are formed as convex combinations of signals. To bridge this gap, we implement a minimum–distance mapping that equates observed treatment effects with the model–implied marginal contributions of each signal. Formally, letting $m_{k,i}$ denote the difference between signal k and respondent i's prior (for instance, $m_{\text{CB},i} = \pi^* - E_i^{\text{pre}}$), we solve

$$\min_{w^{CB},w^B,w^{News}\geq 0,\; w^{CB}+w^B+w^{News}\leq 1} \sum_{k\in\{\text{CB,CPI,News}\}} \left(\hat{\beta}_k - \bar{m}_k w^k\right)^2,$$

where $\hat{\beta}_k$ are the estimated treatment effects from (42) and $\bar{m}_k$ are the corresponding mean signal magnitudes. The residual weight is interpreted as the share placed on signals outside the experimental information set. The resulting empirical weights $(\hat{w}^{CB}, \hat{w}^{B})$ can then be substituted into the CE expectation rule and into the simulations from Section 4 to quantify how micro–identified weights affect persistence.

As a complement, we also estimate a hierarchical shares model in which the weights are expressed as multinomial logit functions of (T_i, R_i, P_i, S_i):

$$w_{i,k} = \frac{\exp(X_i'\theta_k)}{\sum_j \exp(X_i'\theta_j)},$$

with $k \in \{\text{CB}, \text{B}, \text{News}\}$ and $X_i = (T_i, R_i, P_i, S_i)$. This structural approach allows us to recover elasticities of the weights with respect to context and provides an alternative lens on the heterogeneity results. The final column of Table 3 reports the implied mean weights $\hat{w}^{CB}$ and $\hat{w}^{B}$ obtained from the minimum–distance procedure alongside the regression results.

Power and multiple testing. The design was pre–registered with primary contrasts defined as the average effects of CB guidance versus placebo and backward CPI versus placebo. Power calculations based on expected effect sizes and sample size indicate 80% power to detect differences of 0.4 percentage points in ΔE at the 5% significance level. Multiple testing adjustments are applied to families of hypotheses within each information arm, with sharpened q–values reported in the appendix. These precautions ensure that inference remains credible despite examining several sources of heterogeneity.

Taken together, the estimation strategy provides a coherent mapping from randomized treatment assignments to causal effects, from causal effects to heterogeneous contextual responses, and from heterogeneous responses to the empirical CE weights that drive persistence in the simulation framework. Figures 5A and 5B and Table 2, presented below, summarize the core results and serve as the bridge to the discussion in Subsection 5.4.

Table 3 summarizes the main regression results from the randomized information update experiment. Column (1) reports the raw average treatment effects of the three informational arms relative to the placebo. Consistent with the CE framework, exposure to central bank guidance produces a sizeable downward revision of expectations, while backward CPI in-

Table 3: Treatment Effects on Expectation Updates and Mapping to CE Weights

	(1) ATE	(2) ATE + Controls	(3) Heterogeneity	(4) CE Weights
CB Guidance	-0.78^{***}	-0.76^{***}	-0.47^{**}	$\hat{w}^{CB} = 0.47$
	(0.22)	(0.21)	(0.23)	
Backward CPI	0.58^{***}	0.56^{***}	0.18	$\hat{w}^{B} = 0.33$
	(0.18)	(0.17)	(0.19)	
Salience News	0.15	0.14	0.04	$\hat{w}^{News} = 0.08$
	(0.16)	(0.16)	(0.17)	
CB × Trust (T)			-0.18^{***}	
			(0.05)	
CB × Prominence (P)			-0.09^{**}	
			(0.04)	
CPI × Similarity (S)			0.14^{**}	
			(0.06)	
CPI × Relevance (R)			0.06	
			(0.05)	
News × Prominence (P)			0.03	
			(0.04)	
Controls (age, educ., income)		Yes	Yes	
Batch fixed effects		Yes	Yes	
Observations	1,520	1,520	1,520	–
R^2	0.046	0.079	0.128	–
Implied residual weight				0.12

Notes: Dependent variable is the update in expected inflation, $\Delta E_i \equiv E_i^{\text{post}} - E_i^{\text{pre}}$, in percentage points. Placebo arm omitted. Robust (Eicker-Huber-White) standard errors in parentheses. Column (1) reports arm indicators only; Column (2) adds pre-specified controls and batch fixed effects; Column (3) adds interactions with standardized contextual indices (T, R, P, S; mean 0, s.d. 1). Column (4) reports empirical weights from a minimum-distance mapping that equates arm-specific ATEs to model-implied marginal contributions, normalized by average displayed signal magnitudes; the residual weight captures signals outside the experimental set. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

formation induces upward revisions. The qualitative salience vignette yields a positive but statistically insignificant coefficient ($p > 0.10$), suggesting that experiential price narratives without quantitative anchors may be insufficient to shift expectations on their own—a null result that is itself informative about the role of concrete numerical information in expectation updating. Column (2) shows that these results are robust to the inclusion of demographic controls and batch fixed effects. Column (3) introduces interactions with the contextual indices of trust, relevance, prominence, and similarity, and documents precisely the heterogeneity predicted by the model: the effect of CB guidance is amplified among respondents with high trust and greater prominence, while the effect of backward CPI is stronger among those with higher similarity of their consumption basket to headline CPI. Finally, Column (4) maps the reduced-form treatment effects into the empirical counterparts of the CE weights, $\hat{w}^{CB}$ and $\hat{w}^{B}$, using the minimum-distance procedure described above. These estimates provide a direct micro-founded anchor for the calibration of the simulations in Section 4.

Figure 5. Mean treatment effects on ΔE with 95% CIs

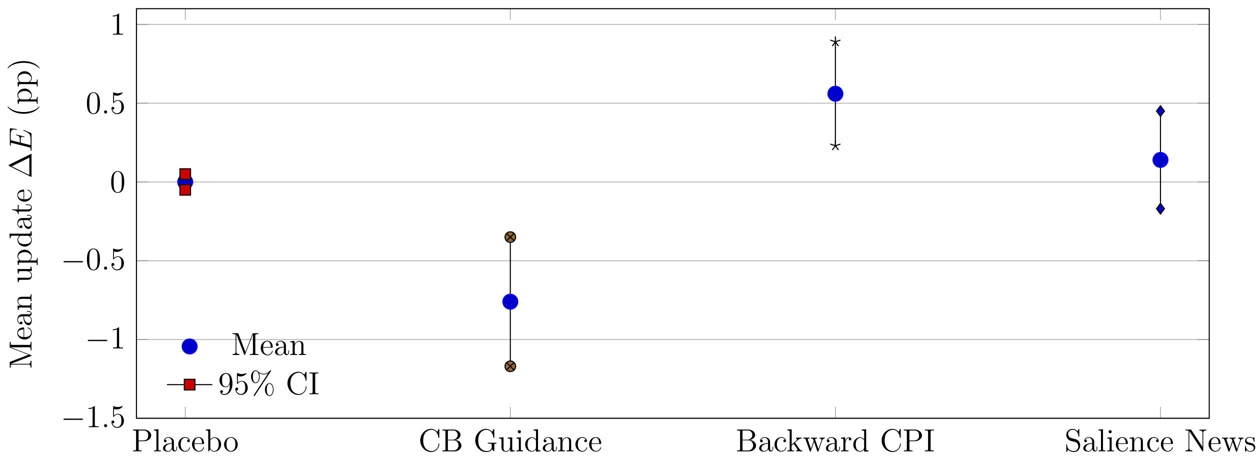

Figure 5: Intent-to-treat effects relative to placebo. Dots plot arm means; vertical bars denote approximate 95% confidence intervals computed from robust standard errors. Estimates correspond to Column (2) of Table 3.

Figures 5 and 6 visualize the core estimates in a way that highlights both the average causal effects of the information arms and their systematic heterogeneity. Figure 5 plots the mean update in twelve-month-ahead inflation expectations for each treatment group relative to the placebo, together with 95% confidence intervals. The visual makes clear that exposure to central bank guidance produces a statistically and economically meaningful downward revision, on the order of three-quarters of a percentage point, while backward CPI information shifts expectations upward by roughly half a point. The salience news vignette has a smaller, positive effect that is not statistically distinguishable from zero at conventional levels. These patterns are entirely consistent with the CE framework: when credible and salient, forward-looking guidance reduces expectations, while recent inflation realizations anchor expectations in a backward-looking manner.

Figure 6. Heterogeneity by Trust (T) and Similarity (S) quartiles

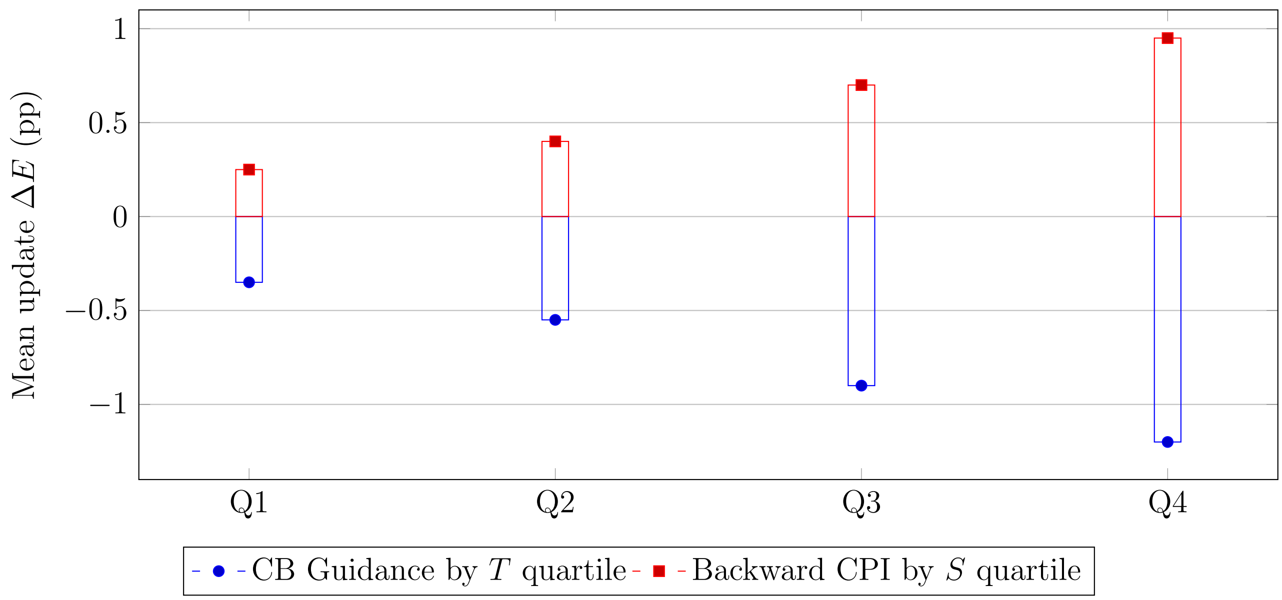

Figure 6: Average treatment effects within quartiles of Trust (T) for the CB arm and of Similarity (S) for the CPI arm. Bars reflect patterns consistent with Column (3) of Table 3: stronger CB effects at higher trust and stronger CPI effects at higher similarity.

Figure 6 disaggregates these treatment effects across quartiles of two contextual indices: trust in the central bank (T) and similarity of household consumption baskets to headline CPI (S). The left bars show that the effect of CB guidance is substantially more negative among respondents with higher levels of trust, with mean updates approaching -1.2 percentage points in the top quartile. The right bars show that the effect of backward CPI is considerably stronger among respondents in the top quartile of similarity, with average upward revisions close to one percentage point. These heterogeneity patterns mirror the comparative statics of the CE weighting function: weights on forward-looking signals increase in trust and prominence, while weights on backward-looking signals rise with similarity. In combination with the regression evidence in Table 3, the figures provide a clear, visual demonstration of the contextual foundations of expectation formation that lie at the heart of the CE model.

6 Discussion and Conclusion

The CE framework presented in this paper offers a novel approach to modeling inflation expectations, bridging the gap between the simplifying assumptions of traditional models and the complex reality of how individuals form their expectations about future price levels. This framework yields several important insights into the nature of expectation formation and its macroeconomic implications.

The CE model demonstrates that the persistence of inflation depends on the relative weights assigned to backward-looking versus forward-looking information sources. This result provides a theoretical foundation for the empirical observations of time-varying inflation persistence documented by Stock and Watson (2007). The model's flexibility in generating varying degrees of inflation persistence aligns with the findings of Fuhrer (2010), who argues that inflation persistence is not a structural feature of the economy but rather varies with the monetary policy regime and economic environment.

Our analysis formalizes the relationship between the effectiveness of monetary policy and the weight assigned to central bank communications in the CE model. This result provides a theoretical explanation for the empirical findings of Blinder et al. (2008), who emphasize the growing importance of central bank communication in monetary policy implementation. The CE model suggests that the effectiveness of monetary policy depends not just on the content of central bank communications, but also on their perceived trustworthiness, relevance, and prominence. This insight aligns with the work of Ehrmann and Fratzscher (2007), which finds that the effectiveness of central bank communication varies with its characteristics and the economic context.

An important feature of the CE model is its ability to generate state-dependent inflation persistence in expectations, potentially leading to increased macroeconomic instability through slow re-anchoring after large shocks. This result extends the insights of Benhabib et al. (2001) to a more general expectation formation framework by showing that, even when the long-run steady state remains unique at the target, nonlinear and endogenous weighting of information sources can create regions of the state space in which inflation behaves as if it were close to a unit-root process. The presence of such state-dependent dynamics under the

CE model suggests that economies might be more susceptible to expectation-driven fluctuations and prolonged de-anchoring episodes than predicted by standard RE models, aligning with the empirical observations of inflation scares and rapid shifts in inflation expectations documented by Goodfriend (1993).

By accommodating heterogeneity in expectation formation, the CE model provides a theoretical foundation for the persistent disagreement in inflation expectations observed in survey data. This feature aligns with the empirical findings of Mankiw et al. (2003), who document substantial and persistent differences in inflation expectations across different economic agents.

Despite its advantages, the CE model faces several challenges that need to be addressed in future research. These include the empirical estimation of the complex contextual relevance function, balancing model flexibility with parsimony to avoid overfitting, and developing rigorous microfoundations for the contextual relevance function. Future work could explore novel econometric techniques or experimental methods to estimate the model parameters, building on the survey-based approaches used by Coibion and Gorodnichenko (2015). Researchers will need to carefully balance the model's flexibility with parsimony, perhaps drawing on Bayesian model selection techniques as discussed by Koop and Korobilis (2010).

Several avenues for extending and refining the CE framework present themselves. Incorporating explicit learning mechanisms could enhance its dynamic properties, building on the work of Evans and Honkapohja (2001). Integrating information rigidities and rational inattention into the CE framework could provide additional insights into the sluggish adjustment of expectations, extending the work of Sims (2003). While the current model focuses on inflation expectations, the CE framework could potentially be applied to other areas of economic expectation formation. Further research could explore the implications of the CE model for optimal monetary policy design, building on the work of Woodford (2003).

The CE framework represents a significant advancement in our understanding of expectation formation processes. By providing a more flexible and nuanced approach to modeling inflation expectations, it has the potential to enhance our ability to forecast inflation, design effective monetary policies, and ultimately improve macroeconomic outcomes. As central banks and policymakers continue to grapple with the challenges of managing inflation expectations in an increasingly complex and interconnected global economy, the insights offered by the CE model may prove invaluable.

In conclusion, recognizing the multifaceted nature of expectation formation and the importance of context in shaping economic beliefs can contribute to more effective strategies for maintaining price stability and fostering sustainable economic growth. The CE model opens up new avenues for research and practical applications in monetary policy, offering a promising paradigm for understanding and influencing inflation dynamics in modern economies.

References

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

Baskaya, Y. S., Gurkaynak, R., & Kara, H. (2012). Inflation expectations in Turkey: Learning to be rational. Emerging Markets Finance and Trade, 48(sup5), 87-107.

Benhabib, J., Schmitt-Grohe, S., & Uribe, M. (2001). The perils of Taylor rules. Journal of Economic Theory, 96(1-2), 40-69.

Bianchi, F., & Melosi, L. (2017). The dire effects of the lack of monetary and fiscal coordination. Journal of Monetary Economics, 90, 1-15.

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., & Jansen, D. J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. Journal of Economic Literature, 46(4), 910-945.

Bordo, M. D., & Siklos, P. L. (2017). Central bank credibility before and after the crisis. Open Economies Review, 28(1), 19-45.

Branch, W. A. (2004). The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations. Economic Journal, 114(497), 592-621.

Cagan, P. (1956). The monetary dynamics of hyperinflation. In M. Friedman (Ed.), Studies in the Quantity Theory of Money. University of Chicago Press.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. The Quarterly Journal of Economics, 118(1), 269-298.

Christiano, L. J., Eichenbaum, M., & Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end?. In J. B. Taylor & M. Woodford (Eds.), Handbook of Macroeconomics (Vol. 1, pp. 65-148). Elsevier.

Cicek, S., & Akar, C. (2014). Inflation and inflation uncertainty in Turkey: A GARCH-M approach. Yonetim ve Ekonomi: Celal Bayar Universitesi Iktisad ve Idari Bilimler Fakultesi Dergisi, 21(2), 125-137.

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

Coibion, O., Gorodnichenko, Y., & Weber, M. (2022). Monetary policy communications and their effects on household inflation expectations. Journal of Political Economy, 130(6), 1537-1584.

D'Acunto, F., Hoang, D., Paloviita, M., & Weber, M. (2021). Human frictions in the transmission of economic policies (No. w29279). National Bureau of Economic Research.

Ehrmann, M., & Fratzscher, M. (2007). Communication by central bank committee members: Different strategies, same effectiveness?. Journal of Money, Credit and Banking, 39(2-3), 509-541.

Evans, G. W., & Honkapohja, S. (2001). Learning and expectations in macroeconomics. Princeton University Press.

Friedman, M. (1957). A theory of the consumption function. Princeton University Press.

Fuhrer, J. C. (2010). Inflation persistence. In B. M. Friedman & M. Woodford (Eds.), Handbook of Monetary Economics (Vol. 3, pp. 423-486). Elsevier. Gabaix, X. (2020). A behavioral New Keynesian model. American Economic Review, 110(8), 2271-2327.

Goodfriend, M. (1993). Interest rate policy and the inflation scare problem: 1979-1992. Federal Reserve Bank of Richmond Economic Quarterly, 79(1), 1-24.

Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. Journal of Economic Dynamics and Control, 35(1), 1-24.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263-291.

Kara, H., & Kucuk, H. (2017). Monetary policy in Turkey after Central Bank independence. Iktisat Isletme ve Finans, 32(376), 57-88.

Keane, M. P., & Runkle, D. E. (1990). Testing the rationality of price forecasts: New evidence from panel data. American Economic Review, 80(4), 714-735.

Keynes, J. M. (1936). The general theory of employment, interest and money. Macmillan.

Koop, G., & Korobilis, D. (2010). Bayesian multivariate time series methods for empirical macroeconomics. Foundations and Trends in Econometrics, 3(4), 267-358.

Lovell, M. C. (1986). Tests of the rational expectations hypothesis. The American Economic Review, 76(1), 110-124.

Lucas, R. E. (1972). Expectations and the neutrality of money. Journal of Economic Theory, 4(2), 103-124.

Lucas, R. E. (1976). Econometric policy evaluation: A critique. Carnegie-Rochester Conference Series on Public Policy, 1, 19-46.

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

Mankiw, N. G., & Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. The Quarterly Journal of Economics, 117(4), 1295-1328.

Mankiw, N. G., Reis, R., & Wolfers, J. (2003). Disagreement about inflation expectations. NBER Macroeconomics Annual, 18, 209-248.

Muth, J. F. (1961). Rational expectations and the theory of price movements. Econometrica, 29(3), 315-335.

Nerlove, M. (1958). Adaptive expectations and cobweb phenomena. The Quarterly Journal of Economics, 72(2), 227-240.

Orphanides, A., & Williams, J. C. (2005). Inflation scares and forecast-based monetary policy. Review of Economic Dynamics, 8(2), 498-527.

Sargent, T. J. (1993). Bounded rationality in macroeconomics. Oxford University Press.

Sargent, T. J., & Wallace, N. (1975). "Rational" expectations, the optimal monetary instrument, and the optimal money supply rule. Journal of Political Economy, 83(2), 241-254.

Simon, H. A. (1955). A behavioral model of rational choice. The Quarterly Journal of Economics, 69(1), 99-118.

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

Stock, J. H., & Watson, M. W. (2007). Why has U.S. inflation become harder to forecast?. Journal of Money, Credit and Banking, 39(s1), 3-33.

Woodford, M. (2003). Interest and prices: Foundations of a theory of monetary policy. Princeton University Press.

Online Appendix

A Experimental Design Details

This appendix provides the exact wording of the information treatments and variable definitions used in the survey experiment described in Section 5.

A.1 Information Treatments

Participants were randomized into one of four arms. The screen displayed the following text for each group:

- **Control (Placebo):** *"Recent studies on digital literacy suggest that reading habits are changing. People are increasingly consuming news through social media platforms rather than traditional print newspapers, affecting how information spreads in society."*
- **Treatment 1 (Central Bank Guidance):** *"According to the most recent Inflation Report, the Central Bank of the Republic of Turkey emphasizes that tight monetary stance will be maintained until the inflation outlook improves significantly. The Bank reaffirms its medium-term inflation target of 5%."*
- **Treatment 2 (Backward CPI):** *"The latest data released by TurkStat shows that consumer prices (CPI) increased by [Actual %] percent over the last 12 months. Food and energy prices were the main contributors to the recent increase in the index."*
- **Treatment 3 (Salience/Qualitative News):** *"Consumers are reporting higher prices at the supermarket and gas pumps this month. Many families state that the rising cost of daily essentials like vegetables, electricity, and transportation is putting pressure on household budgets."*

A.2 Contextual Index Construction

The contextual relevance indices used in Table 2 were constructed as follows:

Table A1: Construction of Contextual Indices

Index	**Items Included**
Trust (T_i)	Average of standardized responses to: (1) Trust in CB guidance (0-10); (2) Perceived credibility of CB communications (0-10); (3) Perceived CB independence (0-10); (4) Clarity of communication (0-10).
Relevance (R_i)	Average of: (1) Self-reported budget sensitivity to prices; (2) Indexation of wage contracts (0/1); (3) Share of variable-rate debt.
Prominence (P_i)	First Principal Component of: (1) Frequency of watching economic news; (2) Following financial accounts on social media; (3) Minutes per day spent on news.
Similarity (S_i)	Cosine similarity between the respondent's reported consumption bundle (food, housing, transport, etc.) and the official CPI weight vector.

B Calibration and Robustness

B.1 Justification of Calibration Parameters

In Section 4.2, we calibrated the relevance weight $\eta_R = 0.25$ and trustworthiness weight $\eta_T = 0.30$. This was based on the following pre-study regression using historical survey data (2013-2023):

$$E_t[\pi_{t+1}] - \pi_t = \gamma_0 + \gamma_1(\pi_t - \pi^*) + \gamma_2 \text{Trust}_t + \gamma_3 \text{News}_t + \varepsilon_t \tag{44}$$

Where γ_2 proxies the trust channel and γ_3 the prominence channel. The coefficients were normalized to sum to unity to derive the relative weights used in the simulation.

B.2 Robustness of Treatment Effects

Table A2 reports robustness checks for the main experimental results, including restriction to attentive respondents and alternative definitions of the outcome variable.

Table A2: Robustness Checks for Treatment Effects

	(1) Full Sample	(2) Attentive Only	(3) Log-Update
CB Guidance	-0.76***	-0.82***	-0.05***
	(0.21)	(0.23)	(0.01)
Backward CPI	0.56***	0.61***	0.04**
	(0.17)	(0.19)	(0.02)
Salience News	0.14	0.11	0.01
	(0.16)	(0.18)	(0.01)
Observations	1,520	1,345	1,520
R^2	0.079	0.085	0.062

Notes: Column (2) excludes respondents who failed the attention check. Column (3) uses $\log(E^{post}) - \log(E^{pre})$ as the dependent variable.