Heterogeneous Inflation Expectations and Subjective Uncertainty: Evidence and Theory

HENG CHEN
University of Hong Kong
YICHENG LIU
University of International Business and Economics
YIBO MENG
University of Hong Kong

October 27, 2026

Abstract. Households' inflation overestimation declines with income, commonly attributed to heterogeneous robust preferences. We propose an alternative: heterogeneous subjective uncertainty. Using the Survey of Consumer Expectations, we document that subjective uncertainty declines with income. We structurally estimate a model in which overconfidence and memory decay shape subjective uncertainty, distinguishing its role from robust preferences in household expectation formation. This distinction also matters for monetary policy: in our estimated economy, overlooking heterogeneous robust preferences has little effect on policy when heterogeneous subjective uncertainty is accounted for, but overlooking the latter has substantial consequences even when preference heterogeneity is recognized.

Keywords. Inflation expectations; subjective uncertainty; robust preferences; overconfidence; memory decay; monetary policy.

JEL Classification. D83, D84, E37

1 Introduction

Households systematically overestimate future inflation (Bhandari et al., 2025), with the magnitude of overestimation declining monotonically with income. This overestimation has consequences at both the household and aggregate levels: it distorts consumption and saving decisions (Vellekoop and Wiederholt, 2019; Coibion et al., 2023) and can amplify business cycles (Pei, 2024). Understanding its sources therefore informs the design of monetary policy. A leading explanation attributes inflation overestimation to robust preferences: households concerned about model misspecification place greater weight on high-inflation outcomes when forming their forecasts (Hansen and Sargent, 2005). In this account, greater overestimation among lower-income households reflects a stronger concern about misspecification (Huo et al., 2024).

This paper explores an alternative source of the income gradient: heterogeneity in subjective uncertainty. Within a framework of robust preferences, the magnitude of overestimation depends jointly on how strongly households guard against model misspecification and how uncertain they perceive future inflation to be. Households with equally strong robust preferences can therefore overestimate inflation by different amounts if their subjective uncertainty differs. A new empirical fact motivates this explanation: households' subjective uncertainty, measured by the dispersion of their belief distributions, declines monotonically with income. The observed decline in overestimation therefore cannot, by itself, establish that the strength of robust preferences declines with income. Our first goal is to incorporate both sources of heterogeneity into a unified framework and use evidence on household expectations to distinguish their respective roles.

Heterogeneous subjective uncertainty is more than a theoretical explanation for the income gradient in overestimation: it also has implications for monetary policy. Differences in subjective uncertainty and robust preferences can generate similar patterns in forecast levels while implying different responses of expectations to policy. A central bank therefore needs to distinguish these two sources of overestimation because optimal policy depends on both what households expect and how their expectations respond to its actions. Understanding the source of heterogeneity also matters for the distribution of policy consequences across income groups. Establishing the policy relevance of heterogeneous subjective uncertainty is the second goal of this paper.

To begin, we document the decline in subjective uncertainty with income using the Survey of Consumer Expectations (SCE), which elicits households' beliefs about inflation over the next 12 months in two ways. Households first report a point forecast, and then assign probabilities to a set of inflation bins. We use the point forecast to measure overestimation, and use a belief distribution to the reported probabilities, taking its standard deviation as our measure of subjective uncertainty. Both moments are con-

structed separately for five income groups, ordered from the lowest to the highest. We find a monotonic pattern: higher-income households are both less pessimistic and less uncertain about future inflation. Furthermore, we complement these findings with evidence on households' overreaction to new information, showing that its magnitude varies little across income groups.

To understand why subjective uncertainty differs across households, we model how households assess and retain information about inflation. A household's subjective uncertainty depends on the precision it assigns to new information and the precision of the beliefs it brings from the past. We introduce two behavioral biases that affect these distinct inputs. Overconfidence leads households to overestimate the precision of new information, reducing their subjective uncertainty. Memory decay makes previously acquired information less precisely recalled, increasing their subjective uncertainty. These biases represent two broader classes of distortions studied in the literature: those affecting beliefs about new information (Bordalo et al., 2020; Afrouzi et al., 2023; Broer and Kohlhas, 2024), and those affecting recalled prior beliefs (Azeredo da Silveira et al., 2024). Incorporating both allows us to study how differences in information processing generate differences in subjective uncertainty and, through robust preferences, inflation overestimation.

In the model, inflation has persistent and transitory components that households do not observe directly. Each household combines a recalled prior with a private noisy signal to form its forecast. Overconfidence distorts the perceived precision of the signal, while memory decay reduces the precision of the recalled prior. Robust preferences then link the resulting subjective uncertainty to overestimation: households concerned about model misspecification place greater weight on adverse, high-inflation scenarios. Greater subjective uncertainty gives these adverse scenarios more scope to influence their forecasts.

We then show that when these two biases interact with robust preferences, they have opposite effects on households' overestimation of inflation. Holding the strength of robust preferences fixed, stronger memory decay raises subjective uncertainty and therefore increases inflation overestimation, whereas stronger overconfidence lowers subjective uncertainty and reduces overestimation. Differences in these information-processing biases can thus generate heterogeneous overestimation even among households with the same strength of robust preferences. Although the two biases have opposite effects on subjective uncertainty and overestimation, both generate overreaction to new information. This difference in their implications helps us distinguish their contributions empirically.

In the next step, we structurally estimate the model to quantify the strength of each feature across the five income groups. For each group, we target subjective uncertainty, the extent of inflation overestimation, and the extent of overreaction. These

three moments together help us identify the strength of each feature. In the data, the extent of overreaction is approximately flat across income groups. Targeting the extent of overreaction therefore helps us pin down the overall extent of the behavioral biases, even though it cannot separate the two. Targeting subjective uncertainty then helps us separate these two biases, since they have direct and opposite implications for subjective uncertainty. Together with the extent of overestimation, this helps us pin down the extent of robust preferences for each income group.

Our estimates reveal systematic differences in all three characteristics across income groups. Across the five groups, estimated overconfidence increases monotonically with income, whereas estimated memory decay decreases monotonically. Both gradients contribute to the decline in subjective uncertainty with income. The estimated strength of robust preferences, however, follows an inverted-U pattern, peaking in the middle-income group and falling toward both ends of the income distribution. Thus, the monotonic decline in overestimation does not reflect a monotonic decline in the strength of robust preferences. It emerges from the joint variation in robust preferences and the subjective uncertainty generated by the two behavioral biases. These estimated profiles are qualitatively consistent with existing behavioral evidence (Belmi et al., 2020; Mani et al., 2013; Mullainathan and Shafir, 2013; Djeutem and Xu, 2025).

Both behavioral biases are necessary in our model to reproduce the joint income profiles of overreaction, subjective uncertainty, and overestimation. Strikingly, the approximately flat profile of overreaction conceals opposing income gradients in the biases that generate it: stronger memory decay contributes to overreaction among lower-income households, while stronger overconfidence contributes to overreaction among higher-income households. Counterfactuals that remove one bias at a time reveal the importance of this offset. Without overconfidence, memory decay alone generates more overreaction among lower-income households and causes the model to overpredict subjective uncertainty and overestimation. Without memory decay, overconfidence alone generates more overreaction among higher-income households and causes the model to underpredict subjective uncertainty and overestimation. Similar degrees of overreaction across income groups therefore mask substantial differences in how households process information.

Monetary policy implications. We ask which source of heterogeneous inflation overestimation matters more for the central bank's policy prescription. We embed the estimated expectation-formation process in an optimal monetary policy problem and reestimate the model under two restrictions on the central bank's assessment of household expectations. In the first, it assumes common robust preferences while allowing subjective uncertainty to vary across income groups. In the second, it assumes common subjective uncertainty while allowing robust preferences to vary. Households retain

their benchmark expectation-formation process in both exercises, so the comparison isolates the consequences of the central bank's assessment for policy and household welfare.

The first restriction has little effect on the policy prescription. When the central bank assumes common robust preferences, the estimated income profiles of overconfidence and memory decay remain close to the benchmark, largely preserving its assessment of how households respond to information. Its chosen inflation rate is very close to the baseline result, approximately 2 basis points above the optimum. The second restriction has much larger consequences. When the central bank assumes common subjective uncertainty, it attributes the declining overestimation profile to declining robust preferences, replacing their inverted-U profile in the benchmark. In doing so, it overlooks differences in the behavioral biases that govern both subjective uncertainty and responses to information. This changes its assessment of both forecast levels and their response to inflation, and its chosen inflation rate falls approximately 19 basis points below the optimum.

The distributional consequences reflect this difference. Under common robust preferences, welfare losses show only limited changes across all income groups. Under common subjective uncertainty, by contrast, the lowest-income group's welfare loss rises by 1.06% of its benchmark loss, while the highest-income group's loss falls by 1.54%. In our estimated economy, allowing robust preferences to differ across income groups does not compensate for overlooking the subjective uncertainty gradient. Heterogeneous subjective uncertainty is therefore an important dimension of household expectations for the central bank's policy choice and its distributional consequences.

Finally, we examine central bank communication as a way of shaping the information households use. We decompose the private signal into a purely private signal and a central bank signal, and quantify the effect of communication by removing the latter. In the benchmark economy, communication lowers the optimal inflation rate by approximately 20 basis points. The effect falls to 7.7 basis points, when subjective uncertainty is homogeneous, and to less than 1 basis point when overconfidence and memory decay are absent. In the benchmark, lower-income households rely more heavily on the central bank signal, so removing communication generates especially large increases in uncertainty and overestimation among households that initially overestimate inflation the most. Overconfidence and memory decay also increase the weight on new information, making subjective uncertainty more responsive to signal precision. Both the heterogeneity in subjective uncertainty and the behavioral biases underlying it are therefore important for understanding the effect of central bank communication.

Related literature. First, our paper contributes to the literature explaining households'

inflation overestimation through robust preferences or ambiguity aversion (Hansen and Sargent, 2008; Bhandari et al., 2025; Michelacci and Paciello, 2024). A central question is whether greater overestimation among lower-income households reflects a stronger preference for robustness. A closely related contribution is Huo et al. (2024), who combine ambiguity aversion with household coordination motives to explain both overreaction and declining overestimation with income. We examine heterogeneous subjective uncertainty as an alternative source of the income gradient within a framework that retains robust preferences. Allowing both sources to vary and using evidence on subjective uncertainty and overreaction helps distinguish their contributions to overestimation. Our estimates show that declining overestimation can coexist with an inverted-U profile of robust preferences, changing the interpretation of the observed income gradient.

Second, we contribute to the empirical literature on heterogeneity in household inflation expectations (Jonung, 1981; Souleles, 2004; Bruine de Bruin et al., 2010; Malmendier and Nagel, 2016; D'Acunto et al., 2021; Kaplan and Schulhofer-Wohl, 2017). This literature documents differences in expectations across demographic groups and examines explanations involving personal inflation experiences, shopping patterns [1], and information. We complement evidence on forecast levels by documenting that households' subjective uncertainty declines monotonically with income. This additional dimension of expectations helps distinguish differences in perceived uncertainty from differences in preferences toward it and provides evidence for interpreting the income gradient in overestimation.

Third, our paper relates to the literature on behavioral biases in expectation formation (Bordalo et al., 2020; Afrouzi et al., 2023; Broer and Kohlhas, 2024; Azeredo da Silveira et al., 2024; Chen and Liu, 2026). These studies examine how departures from Bayesian information processing generate overreaction. We connect such departures to subjective uncertainty by distinguishing distortions in the perceived precision of new information from distortions in recalled prior beliefs. Overconfidence and memory decay both generate overreaction but affect subjective uncertainty in opposite directions. Their opposing income gradients explain how approximately flat overreaction can conceal substantial differences in expectation formation. Subjective uncertainty therefore helps distinguish mechanisms that similar degrees of overreaction alone would leave unresolved.

Finally, we contribute to the literature on optimal monetary policy under incomplete information (Lorenzoni, 2010; Paciello and Wiederholt, 2014; Angeletos and La'O, 2020) and on the distributional consequences of monetary policy under household

[1] Existing literature documents that heterogeneous inflation forecasts across income groups may relate to households' shopping behavior. While the exact source remains debated, possible channels include differences in shopping bundles, shopping channels, and product innovation (Argente and Lee, 2021; Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019).

heterogeneity (Auclert, 2019; Cui and Sterk, 2021). Our contribution is to highlight subjective uncertainty as a policy-relevant dimension of household heterogeneity. We examine how the central bank's recognition of differences in subjective uncertainty affects its policy choice and the distribution of household welfare.

2 Empirical Facts

Using the Survey of Consumer Expectations (SCE), this section documents the empirical patterns that motivate and discipline our analysis. Inflation overestimation and subjective uncertainty both decline monotonically with income. We then examine overreaction, which provides additional evidence on how households process information, and show that its extent is approximately flat across income groups. Together, these patterns cannot be rationalized by a single behavioral bias and help distinguish the sources of heterogeneous household expectations.

2.1 Survey of Consumer Expectations

The SCE is a monthly household survey conducted by the Federal Reserve Bank of New York. Our sample runs from June 2013 to February 2024. Approximately 1,300 households participate each month, and each household remains in the panel for up to 12 months.

The survey elicits beliefs about inflation over the next 12 months in two ways. Households first report a point forecast and then assign probabilities to inflation bins ranging from -12% to 12%. The reported probabilities allow us to approximate each household's subjective distribution of future inflation. Appendix A reproduces the survey questions.

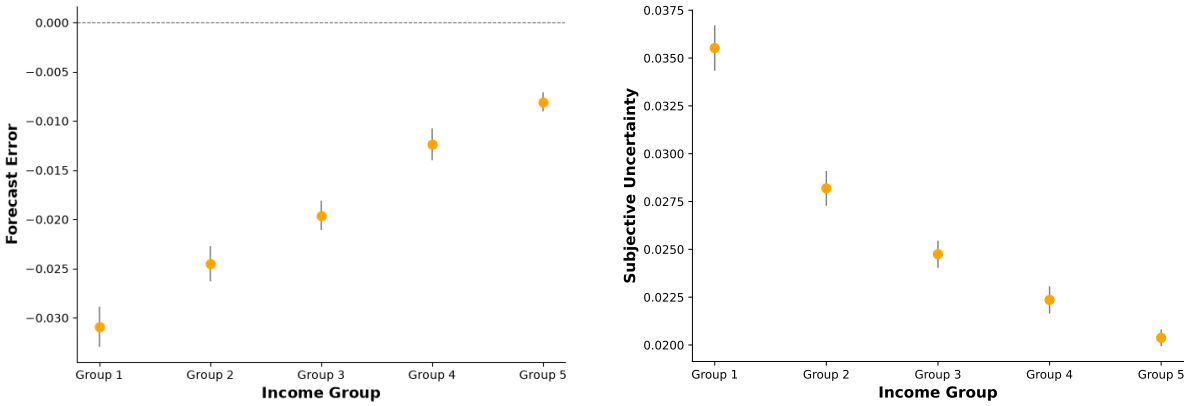
We use point forecasts to measure households' overestimation of the inflation rate, and we use the standard deviation of the fitted belief distribution to measure subjective uncertainty. We construct each measure separately for five income groups, ordered from the lowest- to the highest-income group. Appendix A describes the sample and data construction.

2.2 Overestimation of Inflation across Income Groups

First, we revisit the empirical finding that poorer households tend to overestimate future inflation more than richer households do. Following Huo et al. (2024), we calculate households' forecast error, defined as realized inflation minus the household's point forecast. For household i, the forecast error is

$$FE_{i,t} = \pi_{t+1} - \pi_{i,t+1|t}, \tag{1}$$

where π_{t+1} is realized inflation and $\pi_{i,t+1|t}$ is household i's forecast at time t for inflation over the next 12 months. A negative forecast error then means that the household



(a) Average forecast error

(b) Subjective uncertainty

Figure 1. *Two empirical facts across income groups. Note: Income groups are ordered from the lowest-income group (Group 1) to the highest-income group (Group 5). Panel (a) plots average forecast errors, and Panel (b) plots average subjective uncertainty, measured by the standard deviation of each household's fitted belief distribution. Gray dashed lines denote 95% confidence intervals. Standard errors are clustered at the household level in both panels.*

overestimates the inflation rate.

Panel (a) of Figure 1 shows that average forecast errors are negative in all five income groups and rise monotonically toward zero as income increases. Thus, every group exhibits overestimation, but its magnitude declines monotonically across income groups.

In Appendix B, we provide several robustness checks. First, since households' expectations during the COVID period are changing rapidly, we exclude the COVID period and use the remaining sample to check households' overestimation. Second, instead of using point forecasts, we also use households' belief distributions to calculate the mean forecast and use this mean as our measure of overestimation. We show that all the robustness checks display the same pattern as the baseline: poorer households overestimate the inflation rate more than richer households.

2.3 Subjective Uncertainty across Income Groups

Next, we present a new empirical finding: households' subjective uncertainty, which we define as the dispersion of each household's belief distribution, also declines monotonically across income groups. We use the standard deviation of the belief distribution as our baseline measure of subjective uncertainty, so that a more dispersed distribution represents greater uncertainty and a more concentrated distribution represents lower uncertainty.

For each household, the SCE fits a Beta distribution to the reported bin probabilities and records its variance. We take the square root of this variance to obtain the standard deviation.

Panel (b) of Figure 1 shows that average subjective uncertainty declines monotonically from the lowest- to the highest-income group. The point estimates form a clear income gradient: subjective uncertainty in the lowest-income group is almost twice that in the highest-income group.

Appendix B reports three robustness checks. First, to rule out the influence of the COVID pandemic, we exclude the COVID period and find that overestimation and subjective uncertainty decline across income groups just as in the baseline sample. Second, to address possible measurement error in our baseline construction, we recompute subjective uncertainty using two alternative measures. One follows Engelberg et al. (2009) and fits a continuous belief distribution to each household's reported bin probabilities instead of assigning an ad hoc midpoint to each bin; we describe this construction in detail in the appendix. The other uses the interquartile range, the difference between the 75th and 25th percentiles of the belief distribution, which is less sensitive to the tails of the distribution. In all three checks, the declining pattern of subjective uncertainty across income groups remains qualitatively unchanged.

The documented decline in subjective uncertainty offers an alternative source of the income gradient in inflation overestimation: differences in overestimation may reflect differences in subjective uncertainty as well as differences in the strength of robust preferences. To distinguish these sources, we also examine how households process information. The next subsection introduces an additional moment that captures this dimension of expectation formation.

2.4 Overreaction across Income Groups

The income gradients in overestimation and subjective uncertainty describe the level of households' forecasts and the precision of their beliefs. They do not, by themselves, reveal how strongly households respond to information. This additional dimension matters because different ways of processing information can produce similar responses while having different implications for subjective uncertainty. We therefore examine overreaction across the same five income groups. Together with subjective uncertainty and overestimation, this moment will help distinguish the mechanisms underlying household expectations in the model and estimation that follow.

Overreaction refers to households placing excessive weight on new information relative to a Bayesian benchmark. Following Mincer and Zarnowitz (1969) and Farmer et al. (2024), we measure it using the Mincer-Zarnowitz regression. For each income group, we estimate

$$\pi_{t+1} = \alpha + \beta_{MZ}\pi_{i,t+1|t} + \epsilon_{i,t}. \tag{2}$$

Under rational expectations and the standard orthogonality conditions, β_{MZ} equals one. A coefficient below one indicates that forecasts vary too much relative to subsequent inflation outcomes and is interpreted here as overreaction; a coefficient above

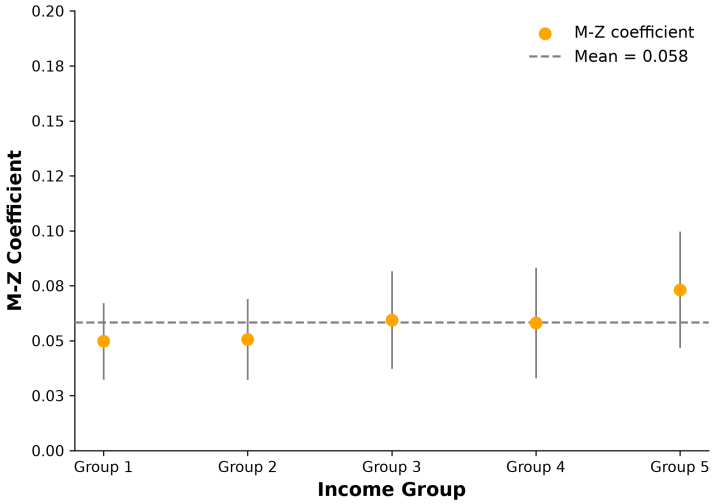


Figure 2. *Mincer-Zarnowitz coefficients across income groups. Note: The figure plots the estimation results of Equation (2) using the subsample of each income group. The estimated Mincer-Zarnowitz coefficients for all income groups are significantly smaller than 1, and there is no significant difference across income groups. The gray lines correspond to 95% confidence intervals. The horizontal dashed line corresponds to the average of the five point estimates. Standard errors are clustered at the year-month and household levels.*

one indicates underreaction.

Figure 2 shows that the estimated coefficients are significantly below one in all five income groups. The estimates vary little across groups and display an approximately flat income profile. Appendix A.3 reports the detailed estimates. Thus, the pronounced declines in overestimation and subjective uncertainty coexist with similar degrees of overreaction across income groups.

These three patterns jointly discipline an explanation of heterogeneous expectations. The model must account for differences in overestimation and subjective uncertainty without predicting comparably pronounced differences in overreaction. In the next section, we connect subjective uncertainty to how households assess new information and recall prior beliefs. Overconfidence and memory decay affect subjective uncertainty in opposite directions while both generating overreaction, providing a way to distinguish differences in information processing from differences in robust preferences.

3 Model

The model combines two sources of inflation overestimation: subjective uncertainty and robust preferences. Overconfidence and memory decay affect the perceived precision of new information and recalled beliefs, respectively, and thereby determine subjective uncertainty. Robust preferences govern how households respond to this uncertainty when choosing their forecasts. We characterize how these ingredients jointly

determine forecast levels and responsiveness to information. We suppress income-group subscripts in the theoretical analysis; Section 4 allows the strengths of the two biases and robust preferences to differ across income groups.

3.1 Setup

Utility function. The economy is populated by a continuum of households indexed by $i \in [0,1]$. Following Chen et al. (2024) and Huo et al. (2024), the utility of household i depends on both the forecast error and the realized inflation[2]:

$$u(\pi_{i,t+1|t}, \pi_{t+1}) = -\frac{1}{2}(\pi_{i,t+1|t} - \pi_{t+1})^2 - \chi\pi_{t+1}, \text{ with } \chi > 0, \tag{3}$$

where $\pi_{i,t+1|t}$ is household i's period-t forecast of inflation in period $t+1$, and π_{t+1} is realized inflation. The first term captures the household's desire to make an accurate forecast by penalizing squared forecast errors. The second term captures its distaste for high inflation: holding forecast accuracy fixed, utility is lower when realized inflation is higher. The parameter χ governs the strength of this second motive.

Data-generating process. Inflation consists of a persistent component and a transitory component:

$$\pi_t = \pi_t^P + \pi_t^T,$$

The transitory component is i.i.d., $\pi_t^T \sim N(0, \sigma^2)$, while the persistent component follows an AR(1) process:

$$\pi_t^P = \rho\pi_{t-1}^P + (1-\rho)\pi^* + \xi_t, \qquad \xi_t \sim N(0, \sigma_\xi^2),$$

where $\rho \in (0,1)$ governs persistence and π^* is the central bank's constant inflation target. The state-generating process is assumed to be common knowledge among all households.

Information structure. Households do not observe inflation directly. Instead, in each period, household i observes a private noisy signal:

$$s_{it} = \pi_t + e_{it}, \qquad e_{it} \sim N(0, \sigma_e^2),$$

where the signal error e_{it} is i.i.d. across households and over time. After observing s_{it}, household i updates its belief and then makes its forecast.

Behavioral biases. Subjective uncertainty depends on the perceived precision of the information used to form beliefs. We introduce two distortions in these inputs: over-confidence raises the perceived precision of new information, whereas memory decay

[2]Adding a quadratic term in realized inflation yields an estimated coefficient close to zero and leaves the quantitative results essentially unchanged. We therefore retain the linear specification to simplify the exposition and clarify the mechanisms.

reduces the precision of recalled beliefs. These distortions affect both the dispersion of the household's approximating belief distribution and the weights assigned to its prior and signal.

Overconfidence. Overconfidence captures households' tendency to overestimate the quality of their private signals. Following Broer and Kohlhas (2024), we assume that household i perceives the signal error as

$$e_{it} \sim N(0, k\sigma_e^2), \qquad k \in (0, 1].$$

The parameter k governs the extent of overconfidence. When $k = 1$, households perceive the signal variance correctly; when $k < 1$, they underestimate it. A lower k therefore represents stronger overconfidence. It makes the signal appear more informative, reducing subjective uncertainty and increasing the weight on new information when updating beliefs.

Memory decay. Memory decay captures households' inability to recall past beliefs precisely, whether because recall is costly or memory is noisy. Following Neligh (2024), we assume that household i perceives its previous-period belief with error:[3]

$$m_{it} = \pi^P_{i,t|t-1} + \omega_{it}, \qquad \omega_{it} \sim N(0, \sigma_\omega^2),$$

where $\pi^P_{i,t|t-1}$ is household i's period-$(t-1)$ belief about π^P_t, and the memory error ω_{it} is i.i.d. across households and over time. The parameter σ_ω^2 governs the extent of memory decay. When $\sigma_\omega^2 = 0$, households perceive their prior beliefs without error; when $\sigma_\omega^2 > 0$, they perceive them with noise. A larger σ_ω^2 therefore represents stronger memory decay. It makes recalled beliefs less informative, raising subjective uncertainty and reducing the weight on the prior when updating.

Robust preferences. Households distinguish uncertainty within their approximating model from concern that the model itself may be misspecified. The former is determined by their information processing; robust preferences govern their response to the latter. Following Hansen and Sargent (2005) and Huo et al. (2024), they treat the approximating model as a useful but potentially imperfect description of the data-generating process. Households evaluate each forecast against the worst-case model among a set of statistically nearby alternatives:

$$\max_{\pi_{i,t+1|t}} \min_{n_t} \mathbb{E}_t \left[u\left(\pi_{i,t+1|t}, \pi_{t+1}\right) n_t + \frac{1}{\alpha} n_t \ln n_t \right] \quad \text{subject to} \quad n_t > 0, \ \mathbb{E}_t[n_t] = 1, \quad (4)$$

[3]Related models of imperfect memory include Afrouzi et al. (2023), Azeredo da Silveira and Woodford (2019), Azeredo da Silveira et al. (2024), Nagel and Xu (2022), and Neligh (2024). Our specification captures a mechanism shared by approaches in which imperfect recall reduces the precision of prior beliefs: households place less weight on recalled information and more weight on new signals. Alternative specifications that reduce prior precision in this way would preserve this qualitative updating mechanism.

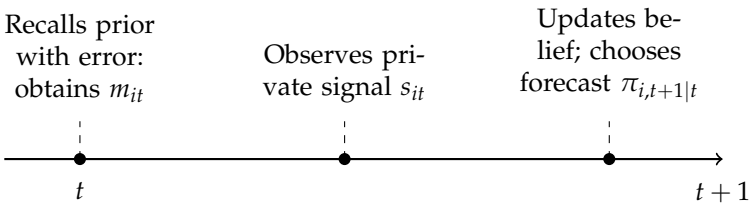


Figure 3. *Timeline of period t.*

where n_t distorts the distribution implied by the approximating model, and $E_t[n_t \ln n_t]$ is the relative entropy of the distorted distribution with respect to that model, scaled by $1/\alpha$. The parameter $\alpha \in (0, \bar{\alpha})$ governs the strength of the household's concern about model misspecification. A larger α lowers the penalty on distortions, allowing nature to move farther from the approximating model and implying a stronger preference for robustness. As $\alpha \to 0$, distortions become infinitely costly and Equation (4) reduces to the standard forecasting problem. The upper bound $\bar{\alpha}$ is the breakdown point beyond which nature can drive utility to $-\infty$ despite the household's optimal response.

Timeline. Figure 3 summarizes the sequence of events in period t. Household i first recalls its prior belief and obtains m_{it}, and then observes the private signal s_{it}. Using both pieces of information, it updates its belief about π_t^P under the approximating model and chooses the forecast $\pi_{i,t+1|t}$ by solving Equation (4).

3.2 Characterization

We first characterize beliefs under the approximating model and then derive the optimal forecast under robust preferences. Throughout, subjective uncertainty refers to dispersion under the approximating model, before the worst-case distortion is applied. This distinction separates the amount of perceived uncertainty from the household's preference for guarding against misspecification. All proofs and closed-form expressions are collected in Appendix D.

Lemma 1. In the steady state, household i's distribution of π_{t+1} under the approximating model has constant variance $\Omega + \sigma^2$, where $\Omega > 0$.

Because the model is linear and Gaussian, the household's filtering variance follows a deterministic recursion. This recursion converges to a unique positive fixed point Ω, so subjective uncertainty about next-period inflation is constant in the steady state.

We impose the following restriction throughout to ensure that the household's max–min problem is well posed.

Assumption 1. $\alpha(\Omega + \sigma^2) < 1$.

Assumption 1 places an upper bound on the strength of robust preferences: $\alpha <$

$1/(\Omega + \sigma^2) \equiv \bar{\alpha}$. Below this breakdown point, the entropy penalty sufficiently restricts the set of distortions considered by nature, preventing worst-case utility from diverging to $-\infty$.

Proposition 1. *Household i's steady-state distribution of π_{t+1} under the approximating model is normal with variance $\Omega + \sigma^2$. The household's optimal forecast under robust preferences is*

$$\pi_{i,t+1|t} = \underbrace{\rho\left[Am_{it} + (1-A)s_{it}\right] + (1-\rho)\pi^*}_{\text{forecast under the approximating model}} + \alpha\chi(\Omega + \sigma^2), \tag{5}$$

where the weight assigned to the recalled prior is $A \equiv \frac{\sigma^2 + k\sigma_e^2}{\Omega + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2} \in (0,1)$.

Equation (5) separates forecasting into two steps. The household first forms a forecast under the approximating model and then adjusts it to guard against misspecification. In the first step, it combines the recalled prior m_{it} and the private signal s_{it} according to their perceived precisions, assigning weights A and $1 - A$, respectively. Overconfidence lowers the perceived variance of the signal, while memory decay raises the perceived variance of the prior. Both biases therefore reduce A and increase the weight placed on the new signal.

In the second step, robust preferences add a strictly positive term $\alpha\chi(\Omega + \sigma^2)$. A malevolent nature places greater weight on low-utility outcomes. Because households dislike high inflation ($\chi > 0$), the resulting worst-case distribution is tilted toward high-inflation states. The household responds by raising its forecast, generating the overestimation of inflation.

Equation (5) exhibits the two sources of heterogeneous overestimation. The upward forecast adjustment depends jointly on subjective uncertainty and the strength of robust preferences: households with identical robust preferences can make different adjustments when their subjective uncertainty differs. The variance $\Omega + \sigma^2$ describes the approximating distribution and is independent of α. Thus, the independence of subjective uncertainty from robustness refers specifically to uncertainty under the approximating model.

3.3 Mechanisms

We derive comparative statics for subjective uncertainty, overestimation, and overreaction. Subjective uncertainty and robust preferences capture the two sources of heterogeneous inflation overestimation: subjective uncertainty measures the dispersion of the approximating distribution, while overestimation measures the upward forecast adjustment added by robust preferences. Overreaction measures the total strength of information-processing distortions, and subjective uncertainty helps separate overconfidence from memory decay. Each result varies one parameter while holding the

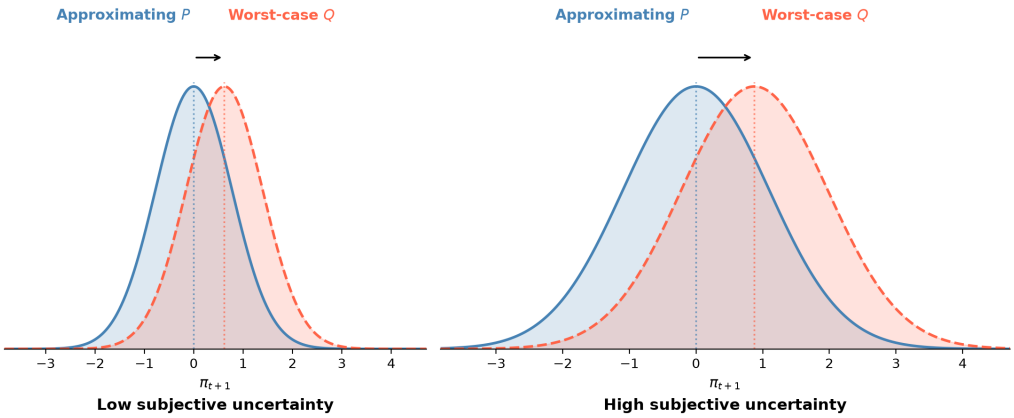


Figure 4. *Worst-case distribution entertained under different levels of subjective uncertainty. Note: This figure compares the worst-case distribution that households with different subjective uncertainty would consider when deciding their optimal forecast. The household on the right panel has larger subjective uncertainty than the household on the left. The blue distribution (P) corresponds to the approximating distribution. The red distribution (Q) corresponds to the worst-case distribution.*

others fixed. All proofs and the conditions used in the comparative statics are collected in Appendix D.

We first show that overconfidence and memory decay affect the subjective uncertainty in opposite directions.

Proposition 2. *Households' subjective uncertainty, $\Omega + \sigma^2$, (i) decreases as overconfidence becomes stronger and (ii) increases as memory decay becomes stronger. It is (iii) constant with respect to the strength of robust preferences.*

The first two results follow from how each bias changes perceived precision. Stronger overconfidence means a lower k: households perceive their private signal as more precise and form a more concentrated posterior, so subjective uncertainty falls. Stronger memory decay means a larger σ_ω^2: households perceive their recalled prior as less precise and form a more dispersed posterior, so subjective uncertainty rises.

The third result reflects the distinct role of robust preferences. The parameter α governs how far households guard against misspecification when choosing their forecast, but it does not change the approximating distribution from which subjective uncertainty is determined. Robust preferences therefore shift the optimal forecast without changing subjective uncertainty.

Next, Proposition 3 shows how the magnitude of overestimation depends on behavioral biases and robust preferences.

Proposition 3. *When $\alpha > 0$, households systematically overestimate future inflation. The overestimation (i) decreases as overconfidence becomes stronger and (ii) increases as memory decay becomes stronger. It also (iii) increases with the strength of robust preferences.*

The first two results operate through subjective uncertainty. Holding the strength of robust preferences (α) and the distaste for inflation (χ) fixed, stronger overconfidence makes new information appear more precise and reduces subjective uncertainty ($\Omega + \sigma^2$). With less perceived uncertainty, adverse high-inflation outcomes have less scope to raise the household's forecast, so overestimation falls. Stronger memory decay instead makes recalled beliefs less precise and increases subjective uncertainty. This gives adverse high-inflation outcomes greater scope to influence the forecast, increasing overestimation.

The third result reflects the direct role of robust preferences. A larger α weakens the entropy penalty and allows nature to consider distributions farther from the approximating model. Because high inflation lowers household utility, nature shifts probability toward high-inflation outcomes, and the household responds by raising its forecast. Stronger robust preferences therefore generate a larger overestimation.

Figure 4 illustrates why subjective uncertainty amplifies the effect of robust preferences. The two households have the same mean belief and the same robust preferences α, but the household in the right panel has a more dispersed approximating distribution P. For a given entropy penalty, nature can shift this distribution farther toward high inflation, producing a larger difference between the approximating distribution and the adverse distribution Q. The household with greater subjective uncertainty therefore exhibits a larger overestimation.

Finally, we examine overreaction, the tendency to place excessive weight on new information relative to the Bayesian benchmark. Section 2.4 introduces the empirical measure, the Mincer-Zarnowitz coefficient β_{MZ} in Equation (2); a lower coefficient indicates greater overreaction. Proposition 4 characterizes how this coefficient varies with memory decay, overconfidence, and robust preferences.

Proposition 4. *The Mincer-Zarnowitz coefficient (i) decreases with stronger memory decay and (ii) decreases with stronger overconfidence. It (iii) remains constant with respect to the strength of robust preferences.*

The first two results reflect a common mechanism. Stronger memory decay makes the recalled prior less precise, while stronger overconfidence makes the private signal appear more precise. Either bias leads households to place less weight on the prior and more weight on the signal. Because the signal contains transitory inflation and noise, forecasts become excessively responsive to information that does not predict future inflation. Households therefore overreact more, and the Mincer-Zarnowitz coefficient falls.

Robust preferences only add a constant upward adjustment to the forecast in Equation (5), leaving its response to information and hence β_{MZ} unchanged.

These comparative statics distinguish the roles of information processing and pref-

erences. Overestimation reflects both subjective uncertainty and robust preferences. Overreaction disciplines the strength of information-processing distortions, while subjective uncertainty helps distinguish their composition because overconfidence and memory decay affect it in opposite directions. These distinctions matter for policy because the same determinants of subjective uncertainty shape both forecast levels and responses to information; Section 5 examines their policy implications.

3.4 Special Cases

We consider two restrictions of the model. First, we remove robust preferences to isolate the effects of overconfidence and memory decay on the updating rule. Second, we retain robust preferences but allow only one behavioral bias at a time. These cases clarify which implications follow from the individual ingredients and which depend on their interaction.

Behavioral biases without robust preferences. First, suppose households exhibit both behavioral biases but have no preference for robustness: $k < 1$, $\sigma_\omega^2 > 0$, and $\alpha \to 0$. The optimal forecast is

$$\pi_{i,t+1|t}^{ov} = \rho\left[Am_{it} + (1-A)s_{it}\right] + (1-\rho)\pi^*. \tag{6}$$

The two biases raise the weight on the private signal above its Bayesian benchmark and therefore generate overreaction. Without robust preferences, however, the additive distortion in Equation (5) vanishes, so average forecast errors are zero. Behavioral biases alone therefore cannot account for the positive overestimation observed in every income group in Figure 1(a).

Both biases affect the weight A assigned to the recalled prior. The updating weight alone therefore need not distinguish overconfidence from memory decay. The following corollary establishes an equivalence in this weight, rather than in the full distribution of forecasts.

Corollary 1. *Without robust preferences, overconfidence and memory decay are observationally equivalent in the weight assigned to new information. For every $k^* \in (0,1]$, there exists a unique σ_ω^{*2} such that an overconfident household with perfect recall and a household with memory decay that correctly perceives its signal place the same weight on the private signal, where*

$$\sigma_\omega^{*2} = \frac{(1-k^*)\sigma_\xi^2\sigma_e^2}{\sigma^2 + k^*\sigma_e^2}.$$

Corollary 1 shows that the same weight on new information can arise from either bias. This equivalence does not imply identical forecast distributions: memory decay also introduces recall noise through m_{it}, and the two biases have different implications for subjective uncertainty. Observing the updating weight alone therefore leaves their composition unresolved.

Robust preferences and a single behavioral bias. We next fix robust preferences and the remaining parameters and vary one behavioral bias at a time. Under these restrictions, changes in subjective uncertainty are linked to changes in overreaction. We consider each single-bias case in turn.

Memory decay only. Suppose households have robust preferences and memory decay but are not overconfident, so $\alpha > 0$, $\sigma_\omega^2 > 0$, and $k = 1$. A larger σ_ω^2 makes the recalled prior less precise, leading households to place less weight on the prior and more weight on the private signal. They consequently overreact more, and the Mincer-Zarnowitz coefficient falls. At the same time, the less precise prior raises the posterior variance and hence subjective uncertainty. Along this comparative static, greater overreaction is associated with greater subjective uncertainty.

Overconfidence only. Now suppose households have robust preferences and overconfidence but recall perfectly, so $\alpha > 0$, $k < 1$, and $\sigma_\omega^2 = 0$. A lower k makes the private signal appear more precise, leading households to place more weight on it. They again overreact more, and the Mincer-Zarnowitz coefficient falls. In contrast to memory decay, however, the perceived increase in signal precision makes the posterior distribution more concentrated and therefore reduces subjective uncertainty. Along this comparative static, greater overreaction is associated with lower subjective uncertainty.

Corollary 2. *Holding robust preferences and the remaining parameters fixed, stronger memory decay in the memory-decay-only model generates both greater overreaction and greater subjective uncertainty. Stronger overconfidence in the overconfidence-only model generates greater overreaction but lower subjective uncertainty.*

Allowing both behavioral biases to vary across income groups helps reconcile declining subjective uncertainty with approximately flat overreaction. Stronger overconfidence and weaker memory decay both reduce subjective uncertainty, but their effects on overreaction work in opposite directions and can offset. The preceding corollary establishes these relationships holding robust preferences and other parameters fixed. Section 4 jointly estimates the mechanisms and evaluates their contributions.

4 Estimation

We structurally estimate the model using empirical moments of inflation overestimation, subjective uncertainty, and overreaction across income groups. The estimates recover heterogeneity in information processing and robust preferences, providing the basis for the policy analysis in Section 5. We first describe the identification strategy and estimation method, then interpret the estimated parameter profiles across income groups and quantify the contributions of the two behavioral biases.

Table 1. *Relationships between group-specific parameters and moments*

Moments	$k \downarrow$	$\sigma_\omega^2 \uparrow$	$\sigma_e^2 \uparrow$	$\alpha\chi \uparrow$
Mincer-Zarnowitz coef.	$\downarrow$	$\downarrow$	$\downarrow$	–
Subjective uncertainty	$\downarrow$	$\uparrow$	$\uparrow$	–
Overestimation	$\downarrow$	$\uparrow$	$\uparrow$	$\uparrow$

Note: This table summarizes how group-specific moments on inflation forecasts vary when the parameters governing behavioral biases and robust preferences vary.

4.1 Identification Strategy

The identification problem is that greater inflation overestimation can reflect greater subjective uncertainty, stronger robust preferences, or both. Subjective uncertainty provides additional information by directly disciplining the dispersion of the approximating distribution. Conditional on this dispersion and the common utility parameters, overestimation helps identify the strength of robust preferences. We estimate the robustness term $\alpha\chi$ rather than α and χ separately. Because χ is common across income groups, cross-group variation in $\alpha\chi$ reflects variation in robust preferences.

Overreaction helps recover the information-processing biases underlying subjective uncertainty. As shown in Section 3.3, overconfidence and memory decay both increase overreaction but have opposite effects on subjective uncertainty. Holding other parameters fixed, lower subjective uncertainty for a given degree of overreaction indicates a larger role for overconfidence, whereas higher subjective uncertainty indicates a larger role for memory decay. These relationships motivate using all three group-specific moments reported in Section 2. Table 1 summarizes the relevant comparative statics. They describe joint identifying relationships, and all parameters are estimated jointly.

We allow overconfidence, memory decay, and robust preferences to differ across income groups, while imposing common utility parameters and a common objective information structure. Income groups index household types; we do not impose a monotonic relationship between income and any of the three group-specific parameters.

A full-sample moment helps discipline the common parameters: the variance of forecast errors informs the variance of private-signal noise, which affects both subjective uncertainty and overreaction.

4.2 Estimation Method

We first calibrate the inflation process using the quarterly personal consumption expenditures (PCE) price index from 2013Q3 to 2023Q4, with inflation measured as the annualized log change in the index. Appendix Table A.2 reports the resulting values of ρ, σ_ξ, σ, and π^*.

We estimate 16 parameters using 16 moments. For each income group g, we estimate overconfidence k_g, memory decay $\sigma^2_{\omega,g}$, and the robustness term $(\alpha\chi)_g$. The signal-error variance σ^2_ϵ is restricted to be common across groups. On the data side, we use the Mincer-Zarnowitz coefficient, subjective uncertainty as measured by the standard deviation of each household's fitted belief distribution, and the overestimation of inflation for each of the five income groups, together with the variance of forecast errors for the full sample. Following Bordalo et al. (2020), we use the full-sample variance of forecast errors to discipline the variance of signal noise. All parameters are estimated jointly.

Following Chernozhukov and Hong (2003), we compute a Laplace-type estimator using Markov chain Monte Carlo. The estimator minimizes a GMM-style quadratic distance between the empirical and model-implied moments:

$$\min_{\theta} \left(M_{data} - M_{model}(\theta)\right)' W \left(M_{data} - M_{model}(\theta)\right), \tag{7}$$

where M_{data} and $M_{model}(\theta)$ collect the empirical and model-implied moments, θ is the parameter vector, and W is a diagonal weighting matrix with entries equal to the inverse sampling variances of the empirical moments.

4.3 Estimation Results and Interpretation

The model closely reproduces the three targeted cross-group profiles: Mincer-Zarnowitz coefficients are approximately flat and below one, while subjective uncertainty and overestimation decline with income. Because these moments enter estimation, this fit describes how the estimated model accounts for the data patterns. Table A.3 reports all parameter estimates, including the common parameters.

The estimated behavioral biases explain how subjective uncertainty declines with income. Overconfidence increases monotonically: higher-income groups have lower k_g, indicating that they perceive their private signals to be more precise. Memory decay decreases monotonically, with lower-income groups having larger $\sigma^2_{\omega,g}$. Both gradients lower subjective uncertainty as income rises. Their implications for overreaction, however, work in opposite directions: stronger overconfidence raises overreaction at higher incomes, whereas weaker memory decay reduces it. These opposing contributions help reconcile declining subjective uncertainty with approximately flat overreaction.

The intriguing finding is that the monotonic decline in inflation overestimation does not translate into a monotonic decline in estimated robust preferences. Figure 5 and Table A.3 show that the robustness term, $(\alpha\chi)_g$, follows an inverted-U profile and is largest for the middle-income group. Once subjective uncertainty is taken into account, greater overestimation among lower-income households need not im-

ply a stronger preference for robustness. The observed income gradient reflects the joint variation in subjective uncertainty and robust preferences. In other words, the joint moments imply systematic heterogeneity in information processing and a non-monotonic profile of robust preferences.

The model treats overconfidence, memory decay, and the robustness term as exogenous and therefore does not explain why their profiles covary with income. Existing behavioral evidence offers possible interpretations of these profiles. Belmi et al. (2020) document greater overconfidence among individuals of higher social class and attribute the relationship to status incentives rather than differences in ability. A stronger incentive to appear competent can bias self-assessments upward, providing a potential interpretation of the increase in overconfidence across income groups in our estimates.

Mani et al. (2013) show that financial scarcity impairs performance on cognitive tasks involving memory and attention. Mullainathan and Shafir (2013) interpret this evidence as a bandwidth effect: financial concerns absorb cognitive resources that would otherwise be available for encoding and retrieving information. This mechanism is consistent with stronger memory decay among lower-income households.

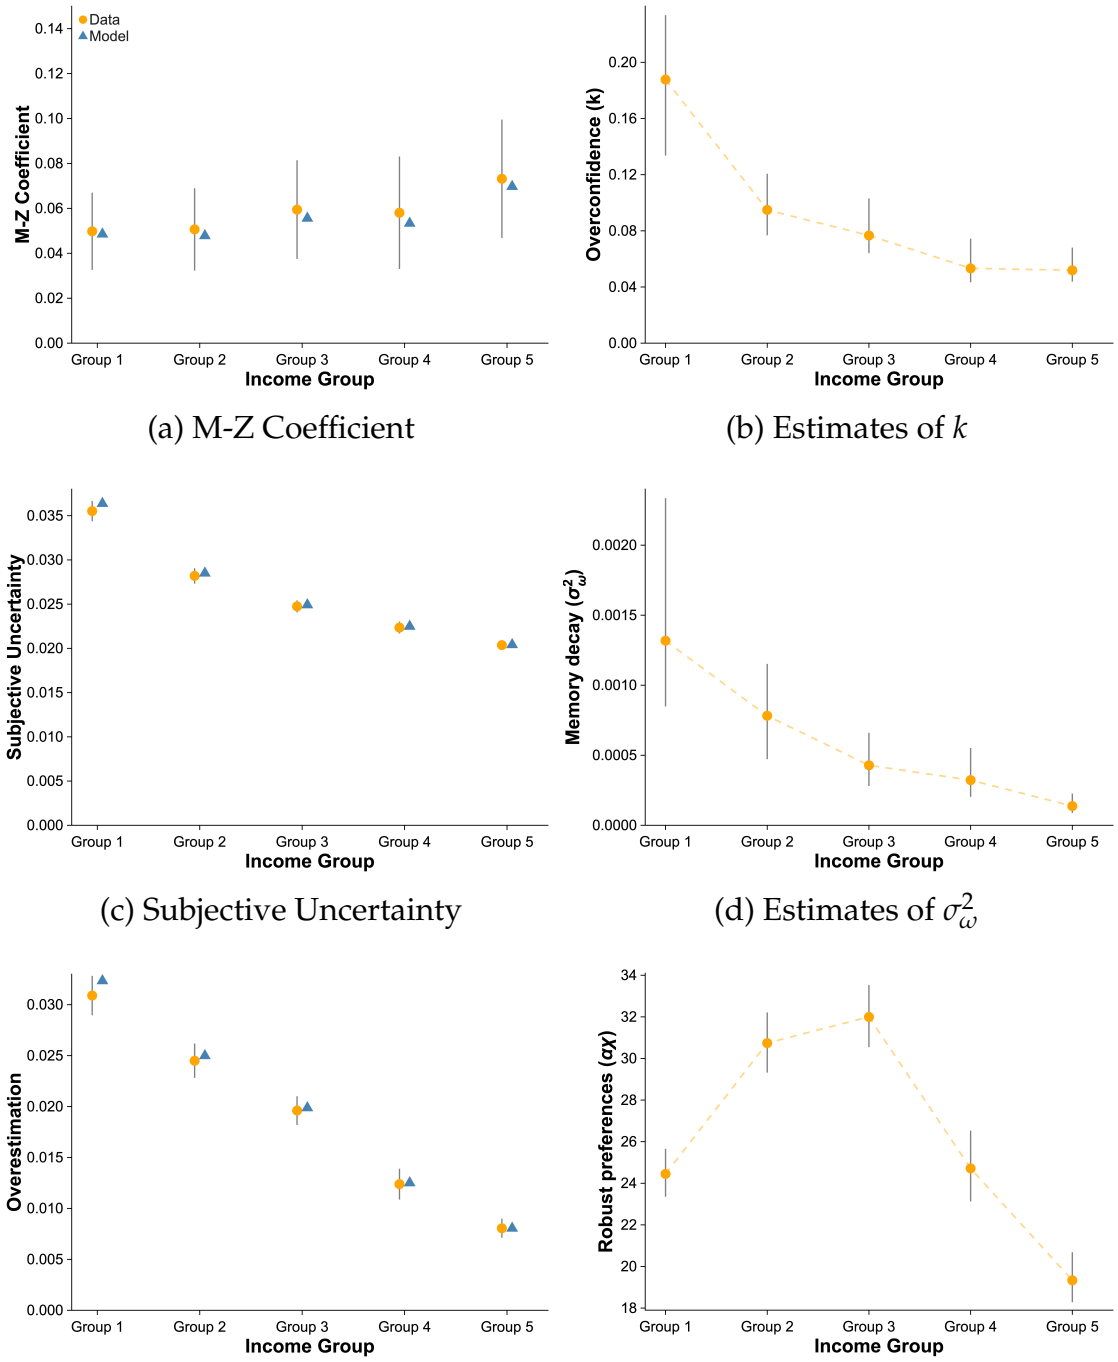
Finally, Djeutem and Xu (2025) show that model distortions under Knightian uncertainty can vary non-monotonically with wealth: distortions are smallest at the bottom and top of the wealth distribution and largest in the middle. Although their application concerns risky-asset returns, this pattern is qualitatively consistent with the inverted-U profile of $\alpha\chi$ in our estimates.

4.4 Contributions of Overconfidence and Memory Decay

We quantify the contribution of each behavioral bias to the estimated expectation-formation process by removing one bias at a time. In the no-overconfidence counterfactual, we set $k_g = 1$ for every group; in the no-memory-decay counterfactual, we set $\sigma^2_{\omega,g} = 0$. In both exercises, all remaining parameters retain their benchmark estimates; we do not re-estimate the restricted models.[4] Figure 6 reports the resulting moment profiles.

Without overconfidence. The remaining memory decay is strongest among lower-income groups, and the Mincer-Zarnowitz coefficient rises with income in Figure 6(a), disrupting the approximately flat empirical profile. Removing overconfidence also eliminates the force that compresses the posterior, causing the model to overpredict both subjective uncertainty and inflation overestimation in Figures 6(c) and 6(e). The overprediction of both moments is largest for the lowest-income group and declines with income; Group 1's magnitude of inflation overestimation is approximately twice its empirical value.

[4]Whereas Section 3.4 characterizes analytical restrictions on the single-bias cases, these exercises measure each bias's contribution in the estimated economy.



(a) M-Z Coefficient

(b) Estimates of k

(c) Subjective Uncertainty

(d) Estimates of σ^2_ω

(e) Overestimation

(f) Estimates of $\alpha\chi$

Figure 5. *Model fit and parameter estimates across income groups. Note: Panels (a), (c), and (e) report the empirical value and the model's prediction for each targeted moment; the gray lines represent* 1.96 *times the data moments' standard errors. Panels (b), (d), and (f) report the estimated extent of overconfidence (k), memory decay (σ^2_ω), and the robustness term ($\alpha\chi$), with gray lines corresponding to the 68% highest posterior density intervals.*

Without memory decay. The remaining overconfidence is strongest among higher-income groups, and the Mincer-Zarnowitz coefficient falls with income in Figure 6(b), again disrupting the approximately flat empirical profile. Removing memory decay also eliminates the force that disperses the prior, causing the model to underpredict subjective uncertainty and inflation overestimation in Figures 6(d) and 6(f). Overestimation becomes nearly flat and uniformly low. The heterogeneous departures in these two exercises reflect the estimated concentration of memory decay, together with the higher observed subjective uncertainty, in lower-income groups.

Taken together, the exercises show that, at the estimated parameter values, removing either bias disrupts the approximately flat income profile of overreaction, even though subjective uncertainty continues to decline with income. The two biases affect subjective uncertainty in opposite directions, while their opposing income gradients help sustain approximately flat overreaction. Their estimated combination therefore matters for reconciling the income profiles of subjective uncertainty and overreaction.

The behavioral biases underlying subjective uncertainty govern both forecast levels and responses to information. Section 5 asks how the central bank's policy choice and the distribution of household welfare change when it assumes common robust preferences or common subjective uncertainty.

5 The Policy Relevance of Subjective Uncertainty

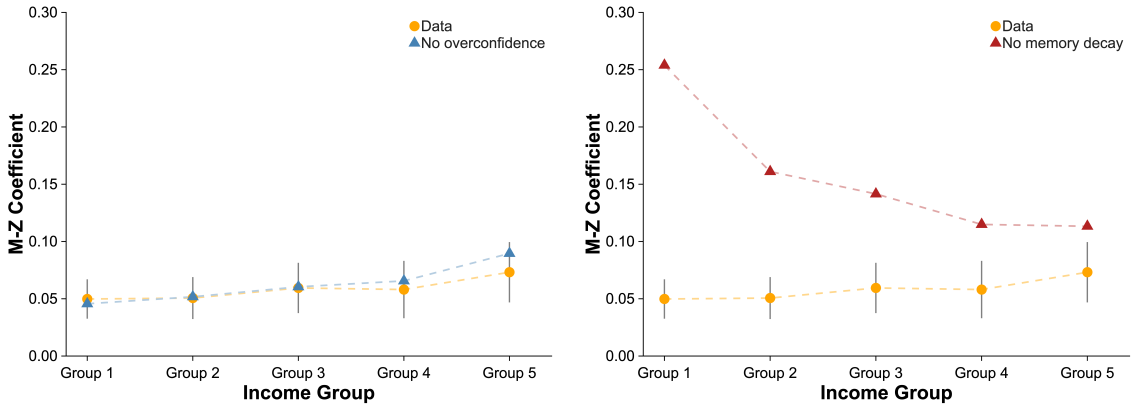
The preceding analysis distinguishes subjective uncertainty and robust preferences as two sources of heterogeneous inflation overestimation across income groups. We now examine the importance of recognizing each source of heterogeneity for the central bank's policy choice.

We embed the estimated expectation-formation process in an optimal monetary policy problem and compare the optimal benchmark under the full estimated model with two counterfactuals. In Counterfactual I, the central bank assumes common robust preferences while allowing subjective uncertainty to vary across income groups. In Counterfactual II, it assumes common subjective uncertainty while allowing robust preferences to vary. Households retain their benchmark expectation-formation process in both exercises; only the central bank's assessment of that process changes. Comparing each counterfactual with the benchmark, in which the central bank recognizes both sources of heterogeneity, reveals the importance of each source for policy choices and household welfare.

5.1 Policy Environment and the Optimal Benchmark

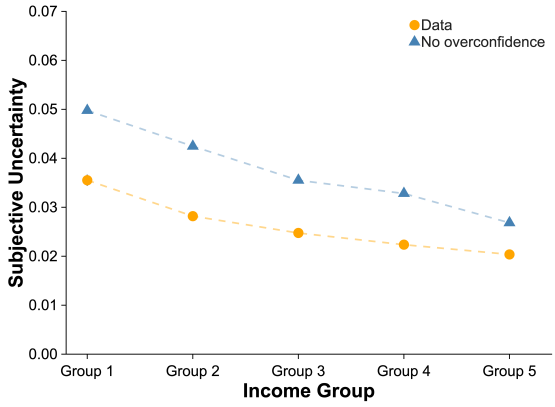
Following Svensson (1997) and Huo et al. (2024), the central bank chooses inflation π to minimize

$$L_{\text{CB}}(\pi) = \frac{1}{2}(\pi - \pi^*)^2 + \frac{\Psi}{2} l^2, \tag{8}$$

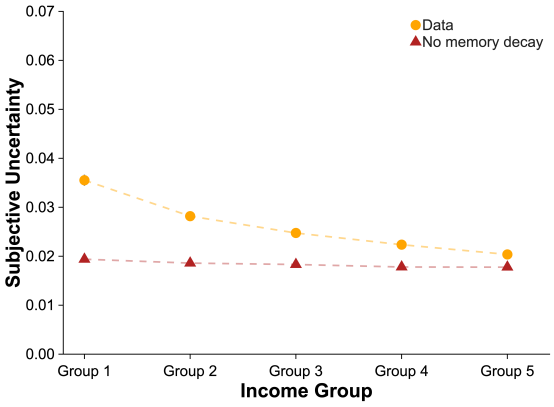


(a) M-Z Coefficient

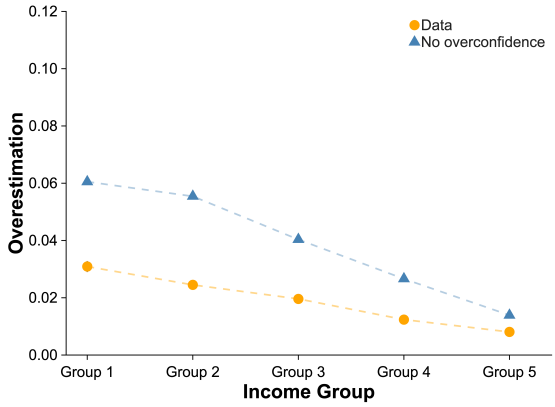
(b) M-Z Coefficient



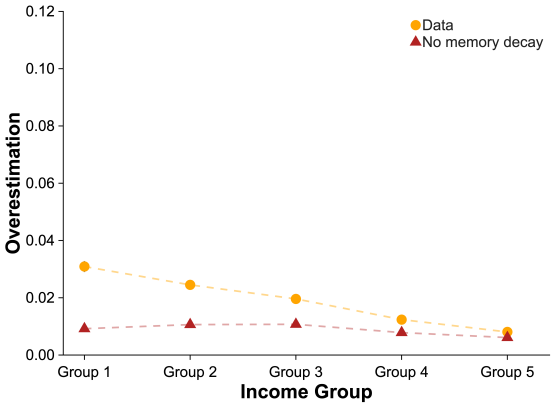
(c) Subjective Uncertainty



(d) Subjective Uncertainty



(e) Overestimation



(f) Overestimation

Figure 6. *Note: This figure reports the empirical value, together with the moments implied by the "no overconfidence" (left column) and "no memory decay" (right column) counterfactual economies. The gray lines correspond to* 1.96 *times the data moments' standard errors.*

where π^* is the inflation target, l is the log ratio of actual to full employment, and $\Psi > 0$ is the relative weight on employment stabilization. The expectations-augmented Phillips curve is

$$l = \theta(\pi - \pi^e),$$

where π^e is the population-average inflation expectation and $\theta > 0$. Policy therefore depends on how expectations respond to inflation as well as on their level.

Households retain the expectation-formation mechanisms in Section 3. For household i in income group g, the recalled prior is $m_{gi} = \pi^* + \omega_{gi}$ and the private signal is $s_i = \pi + e_i$, with variances $\sigma^2_{g,\omega}$ and σ^2_e, respectively. The household perceives the signal-error variance as $k_g\sigma^2_e$. Its forecast in this policy environment is[5]

$$\pi_{gi} = \hat{A}_g m_{gi} + (1 - \hat{A}_g)s_i + (\alpha\chi)_g \Xi_g, \tag{9}$$

where

$$\hat{A}_g = \frac{k_g\sigma^2_e}{\Omega_g + \sigma^2_{g,\omega} + \sigma^2 + k_g\sigma^2_e}, \qquad \Xi_g = \frac{(\Omega_g + \sigma^2_{g,\omega} + \sigma^2)k_g\sigma^2_e}{\Omega_g + \sigma^2_{g,\omega} + \sigma^2 + k_g\sigma^2_e}.$$

The prior variance of persistent inflation, Ω_g, comes from the estimated model. The weight on the recalled prior is $\hat{A}_g$, and Ξ_g is the subjective variance of inflation in this environment. Robust preferences generate the additive adjustment $(\alpha\chi)_g\Xi_g$.

Let r_g denote group g's population share, with $\sum_g r_g = 1$. Averaging forecasts gives

$$\pi^e = \sum_g r_g \pi^e_g = \bar{A}\pi^* + \bar{B}\pi + \bar{C}, \tag{10}$$

where

$$\bar{A} = \sum_g r_g \hat{A}_g, \qquad \bar{B} = \sum_g r_g(1 - \hat{A}_g), \qquad \bar{C} = \sum_g r_g(\alpha\chi)_g\Xi_g.$$

The slope $\bar{B}$ governs how current inflation passes into aggregate expectations, while $\bar{A}\pi^* + \bar{C}$ determines their intercept. A stronger response of expectations to inflation reduces the employment effect of a given increase in inflation. The central bank therefore needs to assess both the level of expectations and their response to inflation when choosing its policy.

Substituting Equation (10) into the Phillips curve and minimizing Equation (8) yields

$$\pi^{\text{opt}} = \frac{1 + \Psi\theta^2(1 - \bar{B})\bar{A}}{1 + \Psi\theta^2(1 - \bar{B})^2}\pi^* + \frac{\Psi\theta^2(1 - \bar{B})}{1 + \Psi\theta^2(1 - \bar{B})^2}\bar{C}. \tag{11}$$

The benchmark central bank constructs these aggregates from the full group-specific

[5]Households do not know the central bank's policy rule. This assumption is consistent with evidence that households are largely inattentive to monetary policy (Coibion et al., 2022; De Fiore et al., 2021).

estimates in Section 4. Throughout the counterfactuals, households retain this same estimated process; only the central bank's model of it changes. Each resulting policy is evaluated in the same underlying economy. Appendix E provides the derivations.

5.2 Two Interpretations of the Income Gradient

The two counterfactuals ask whether the central bank can simplify its account of household heterogeneity without materially changing its policy prescription. We first allow it to overlook differences in robust preferences while retaining differences in subjective uncertainty. We then ask whether it can instead overlook differences in subjective uncertainty while retaining differences in robust preferences. In each case, we reestimate the parameters under the central bank's restriction, allowing its model to account for the targeted evidence as well as possible within that restriction. This reveals how overlooking one source of heterogeneity changes the central bank's interpretation of household expectations and, subsequently, its policy choice. Households retain their benchmark parameters throughout; only the central bank's assessment changes.

Counterfactual I: The central bank assumes common robust preferences. The central bank allows overconfidence and memory decay, and hence subjective uncertainty, to vary across income groups, but assumes a common strength of robust preferences. Differences in overestimation in its model must therefore arise from differences in subjective uncertainty. We estimate this restricted model using the same moments as the benchmark: each group's subjective uncertainty, overestimation, and Mincer-Zarnowitz coefficient, together with the full-sample variance of forecast errors. Appendix Table A.4 reports the targeted moments and their model-implied values.

The left column of Figure 7 shows that this restriction leaves the estimated income profiles of the behavioral biases close to the benchmark. The parameter k_g continues to decline with income, implying stronger overconfidence among higher-income households, while memory decay also declines with income. The central bank therefore retains a similar account of the differences in information processing that generate the subjective uncertainty gradient, even though it replaces the benchmark's inverted-U profile of robust preferences with a common value.

Counterfactual II: The central bank assumes common subjective uncertainty. We next reverse the restriction. The central bank assumes common overconfidence and memory decay, and hence common subjective uncertainty, while allowing robust preferences to vary across income groups. Differences in overestimation in its model must now arise from differences in robust preferences. Because it overlooks the income gradients in the behavioral biases, the estimation targets average subjective uncertainty and the average Mincer-Zarnowitz coefficient rather than their group-specific values. It continues to target each group's overestimation and the full-sample variance of forecast errors.

Appendix Table A.5 reports the targeted moments and their model-implied values.

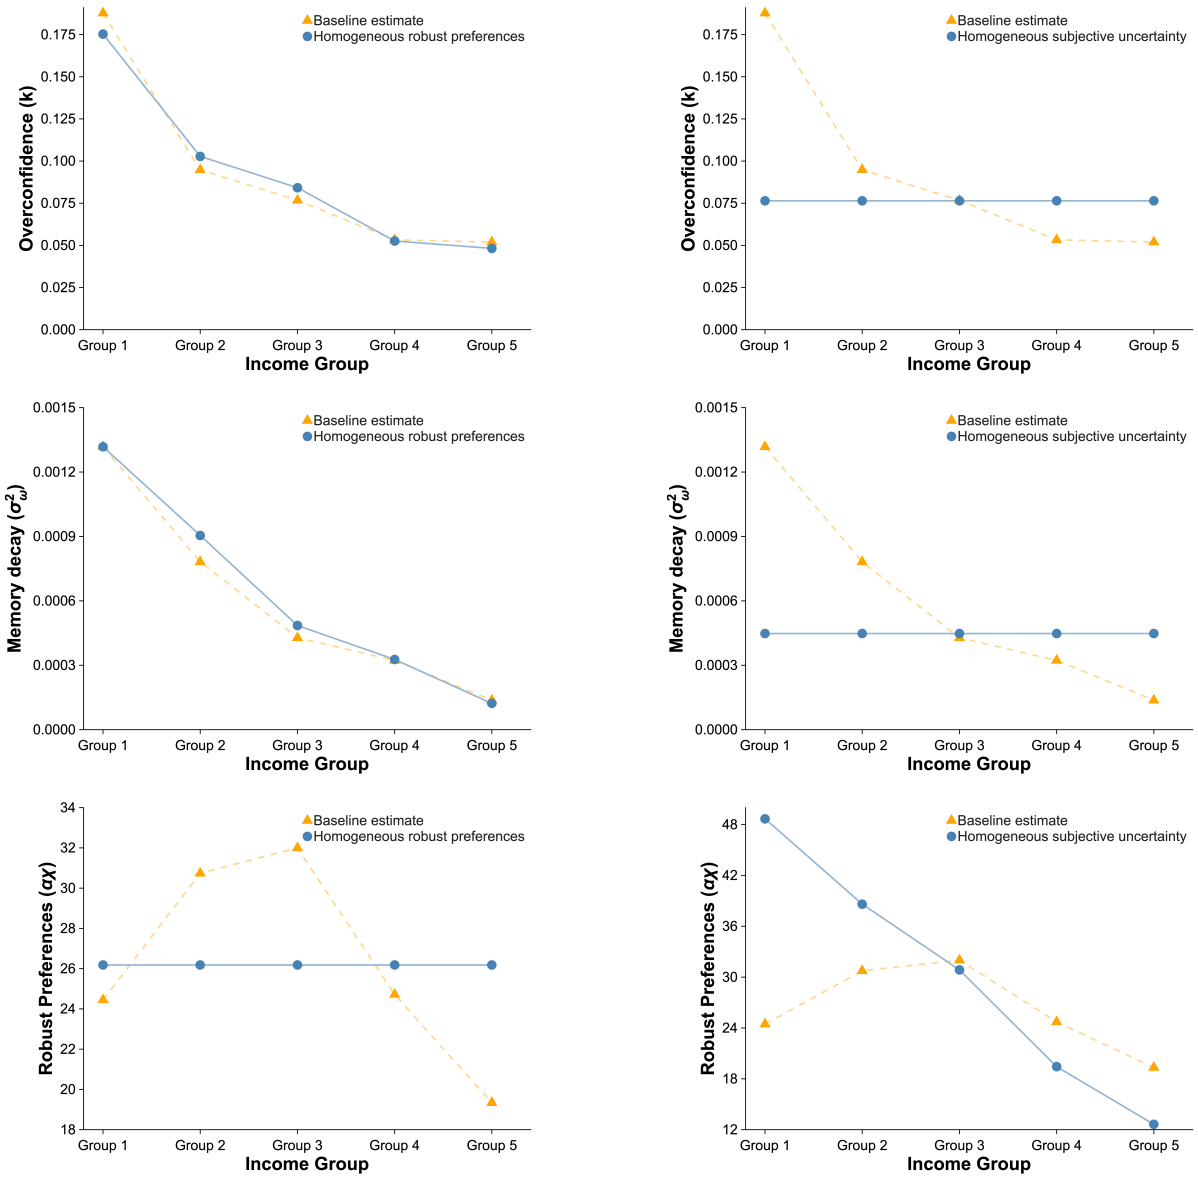
The right column of Figure 7 shows how Counterfactual II changes the central bank's interpretation of the income gradient. With subjective uncertainty held common across groups, the declining overestimation profile is attributed entirely to differences in robust preferences. Common values replace the benchmark income profiles of overconfidence and memory decay, while the estimated robustness term declines monotonically with income instead of following an inverted-U pattern. Relative to the benchmark, this interpretation assigns substantially stronger robust preferences to the lowest-income group and weaker robust preferences to the highest-income group.

The intuition follows from the distinct roles of subjective uncertainty and robust preferences. In Counterfactual I, group-specific uncertainty and overreaction moments continue to discipline the behavioral biases. The estimated bias profiles remain close to the benchmark, largely preserving the central bank's assessment of the weights households assign to information. The restriction mainly changes its assessment of the additive upward forecast adjustment, $(\alpha\chi)_g\Xi_g$. In Counterfactual II, common subjective uncertainty prevents differences in perceived dispersion from explaining the overestimation gradient, so the estimated robustness terms absorb that gradient. At the same time, common overconfidence and memory decay imply common weights on the recalled prior and the current signal. Adjusting robust preferences can change forecast levels, but cannot restore the missing differences in responsiveness to information: in Equation (9), robust preferences enter only the additive adjustment, whereas the response to current inflation depends on $1 - \hat{A}_g$. The two restrictions therefore differ in how they change the central bank's assessment of both forecast levels and responses to inflation. The next subsection quantifies the consequences for policy.

5.3 Consequences for Policy

Table 2 compares the policy choices. The baseline optimal inflation rate is 5.23%. In Counterfactual I, the central bank assumes common robust preferences and chooses 5.25%, approximately 2 basis points above the optimum. In Counterfactual II, it assumes common subjective uncertainty and chooses 5.04%, approximately 19 basis points below (or 3.6% relative to) the optimum. Ignoring heterogeneity in robust preferences thus shifts policy only slightly, whereas ignoring heterogeneity in subjective uncertainty produces a much larger deviation.

To explain this difference in policy outcomes, we examine how each restriction changes the central bank's assessment of forecast levels and the responsiveness of expectations to inflation. When the central bank assumes common robust preferences, it continues to recognize differences in subjective uncertainty and in households' re-



(a) Central bank assumes common robust preferences

(b) Central bank assumes common subjective uncertainty

Figure 7. *Parameter estimates under the central bank's alternative assumptions. Note: In Counterfactual I, the central bank assumes common robust preferences (left column); in Counterfactual II, it assumes common subjective uncertainty (right column). The figure compares the resulting estimates with the benchmark estimates in Section 4. Within each column, the panels report, from top to bottom, overconfidence, memory decay, and robust preferences. Circles denote estimates under the indicated restriction, and triangles denote the benchmark estimates.*

Table 2. *Policy choices under alternative interpretations of heterogeneous overestimation*

Central bank's model	Inflation	Deviation from optimum
Full heterogeneity	5.23%	—
Common robust preferences	5.25%	+2 basis points
Common subjective uncertainty	5.04%	−19 basis points

Note: In Counterfactual I, the central bank assumes common robust preferences and allows subjective uncertainty to vary across groups. In Counterfactual II, it assumes common subjective uncertainty and allows robust preferences to vary across income groups. Households retain the full estimated process in both cases. Policy deviations are approximate and calculated from the displayed inflation rates.

sponses to new information. The estimated overconfidence and memory decay remain close to their benchmark values, so the central bank's assessment of how households process information changes little. The remaining error mainly concerns how strongly each group adjusts its forecast upward in response to a given amount of subjective uncertainty. In our estimates, these errors in robust preferences produce only a small change in the central bank's policy choice. Preserving heterogeneous subjective uncertainty thus retains the features of household expectations that matter most for policy in this exercise.

When the central bank assumes common subjective uncertainty, it underestimates subjective uncertainty among lower-income households and overestimates it among higher-income households. These errors matter for policy because subjective uncertainty affects how much households overestimate inflation: greater subjective uncertainty induces a larger upward forecast adjustment under robust preferences. Moreover, the overconfidence and memory decay that determine subjective uncertainty also determine how much households respond to new information. The central bank's assumption of common values for these biases therefore changes its assessment of both forecast levels and their response to inflation. Both affect the employment consequences of its policy.

For policy design, ignoring heterogeneity in subjective uncertainty generates a substantial policy error relative to the optimum. The central bank's excess loss under each restriction is derived in Appendix E. We next examine how these policy errors affect household welfare across income groups.

5.4 Distributional Consequences

We next examine who bears the consequences of these policy errors. We measure household welfare loss as the negative of utility in Equation (3). For income group g, the average loss is

$$L_g(\pi) = \frac{1}{2}V_g + \frac{1}{2}\left[\pi_g^e(\pi) - \pi\right]^2 + \chi\pi, \tag{12}$$

Table 3. *Household welfare consequences of the two policy errors*

Income group	Additional loss relative to group benchmark Central bank assumes:	
	Common robust preferences	Common subjective uncertainty
Group 1 (lowest income)	-0.06%	1.06%
Group 2	-0.02%	0.42%
Group 3	$\approx 0.00\%$	-0.02%
Group 4	0.03%	-0.47%
Group 5 (highest income)	0.09%	-1.54%

Note: Entries report $100[L_g(\pi^j) - L_g(\pi^{\text{opt}})]/L_g(\pi^{\text{opt}})$, using the household loss in Equation (12). Positive values indicate additional loss; negative values indicate gains. The calculations retain the benchmark group-specific household parameters, equal population shares, $\Psi = 1$, and $\chi = 0.1$. Group 3's change when the central bank assumes common subjective uncertainty rounds to zero. Tables A.6 and A.7 report the underlying components.

where V_g is the within-group variance of forecasts generated by memory and signal errors. Since the true household process is fixed, V_g is independent of the policy choice. Policy changes welfare through the group's average forecast error and the direct disutility from inflation. Appendix E gives the full decomposition, and Tables A.6 and A.7 report its components.

Table 3 reports each group's percentage change in welfare loss relative to the benchmark. In Counterfactual I, every group's loss changes by less than 0.1% of its benchmark value. Counterfactual II produces larger changes, with losses for lower-income households and gains for higher-income households. The lowest-income group's loss falls by 0.06% in Counterfactual I but rises by 1.06% in Counterfactual II; Group 2's loss falls by 0.02% and rises by 0.42%, respectively. The highest-income group's loss instead rises by 0.09% in Counterfactual I and falls by 1.54% in Counterfactual II. These changes measure household welfare; the central bank's objective in Equation (8) is a separate loss measure.

Counterfactual I leaves policy close to the optimum, so household welfare changes little. Counterfactual II lowers inflation by approximately 19 basis points. Lower inflation reduces the direct disutility from inflation for all groups, but also changes their average forecast errors. At the benchmark, lower-income groups forecast inflation above the chosen rate, whereas higher-income groups forecast below it. Under the estimated forecast rules, the lower counterfactual inflation rate increases the average forecast-error loss of lower-income households and reduces that of higher-income households. For lower-income groups, the increase in forecast-error loss outweighs the direct benefit from lower inflation.

These results establish the policy significance of distinguishing the two sources of

heterogeneous overestimation. In our estimated economy, a central bank that recognizes heterogeneous subjective uncertainty but assumes common robust preferences chooses a policy close to the optimum, generating small household welfare changes. A central bank that assumes common subjective uncertainty and attributes the income gradient to heterogeneous robust preferences chooses a policy substantially farther from the optimum, with losses concentrated among lower-income households. Accounting for the subjective uncertainty underlying the observed overestimation gradient is therefore important for both the policy prescription and its distributional consequences.

5.5 Central Bank Communication and Subjective Uncertainty

The analysis so far takes households' information as given and examines how the central bank's interpretation of household expectation formation affects its policy choice. Central bank communication, however, can also shape the information households use. In our model, central bank communication reduces the overestimation of inflation by lowering subjective uncertainty. The resulting reduction in inflation expectation can, in turn, lower the inflation rate chosen by the central bank. We quantify this effect and show how the effect depend on the heterogeneity in subjective uncertainty and the difference between subjective and objective uncertainty.

Decomposing the private signal. We split the signal in Section 5.1 into a purely private signal and a signal from the central bank:

$$x_{gi} = \pi + v_{gi}, \quad v_{gi} \sim N(0, \sigma_v^2), \qquad\qquad z = \pi + u, \quad u \sim N(0, \tau^{-1}).$$

Private noise v_{gi} is independent across households. Communication noise u is common to all households. The parameter τ measures the precision of central bank communication.

Overconfidence now applies to the purely private signal. A household in group g perceives the variance of v_{gi} as $k'_g\sigma_v^2$, with $k'_g \in (0, 1]$, and perceives the precision of the central bank signal correctly. Because both signals are Gaussian, the household's inference depends on them only through their precision-weighted combination $s_{gi} = w_g x_{gi} + (1 - w_g)z$, where $w_g = 1/(1 + \tau k'_g\sigma_v^2)$. The perceived variance of this composite signal is

$$S_g \equiv \frac{\tau^{-1} k'_g\sigma_v^2}{\tau^{-1} + k'_g\sigma_v^2}. \tag{13}$$

We discipline the decomposition by requiring it to reproduce the perceived signal variance estimated in Section 4:

$$\frac{\tau^{-1} k'_g\sigma_v^2}{\tau^{-1} + k'_g\sigma_v^2} = k_g\sigma_e^2, \qquad g = 1, \ldots, 5. \tag{14}$$

Equation (14) reinterprets the estimated information structure without changing it. The updating weights $\hat{A}_g$, subjective uncertainty Ξ_g, and the robust adjustment $(\alpha\chi)_g\Xi_g$ depend on household information only through its perceived variance. They therefore remain the same as in Section 5.1, and the optimal policy is still given by Equation (11). Equation (14) also implies that the weight on the central bank signal within the composite is $1 - w_g = \tau k_g \sigma_e^2$. Households that perceive their information as less precise rely more on the central bank communication.

Calibrating the public signal. Similar to Section 4, τ and σ_v^2 are common across groups, while k'_g varies. Given (τ, σ_v^2), Equation (14) pins down each group's private overconfidence:

$$k'_g = \frac{k_g \sigma_e^2}{\sigma_v^2 \left(1 - \tau k_g \sigma_e^2\right)}.$$

The calibration therefore needs two moments for two unknowns. We use two full-sample moments: the variance of individual forecast errors, already targeted in Section 4, and the variance of monthly mean forecast errors. Averaging forecasts within a month removes private signal noise and memory errors, but not communication noise, which is common across households. The variance of monthly mean forecast errors therefore disciplines τ, and the gap between the two variances disciplines σ_v^2. Appendix C.6 derives both moments and gives each parameter in closed form.

Table A.8 reports the implied signal variances and group-specific parameters. The five k'_g decline with income, preserving the overconfidence gradient from Section 4. Because τ is common across groups, greater overconfidence in private information among higher-income households implies less reliance on central bank communication.

The effect of communication. Given our calibration, to what extent does central bank communication reduce the household subjective uncertainty and therefore the inflation rate chosen by the central bank? How does the effect of central bank communication depend on the heterogeneity in subjective uncertainty and on the difference between subjective and objective uncertainty?

To answer the first question, we remove the communication by letting $\tau \to 0$ to quantify the its effect. The composite signal then reduces to the purely private signal, and each group's perceived signal variance rises from $k_g \sigma_e^2$ to $k'_g \sigma_v^2$. All other parameters stay at their Section 4 values. The central bank still chooses policy by Equation (11), with full knowledge of how households form expectations.

Optimal inflation rises from 5.235% to 5.433%. Central bank communication thus lowers the chosen inflation rate by about 20 basis points. That is 5.8% of the inflation bias, $\pi^{opt} - \pi^*$, that would prevail without communication.[6]

[6]Throughout Section 5, $\theta = 5.6$, computed from Ball and Mazumder (2019).

The effects on policy works through the level of expectations. Communication makes households' information more precise and lowers their subjective uncertainty. This reduction is larger when households place greater weight on new information. Lower subjective uncertainty gives adverse high-inflation scenarios less room to raise forecasts. That shrinks each group's robust adjustment $(\alpha\chi)_g\Xi_g$, and the intercept $\bar{C}$ of aggregate expectations. Higher expected inflation lowers employment at any given inflation rate, so the central bank accommodates part of it, as the second term of Equation (11) shows. By lowering $\bar{C}$, communication therefore lowers the inflation rate chosen by the central bank.

Comparison economies. To understand how the effect of central bank communication depend on the heterogeneity in subjective uncertainty and the gap between subjective and objective uncertainty, we repeat this experiment in two comparison economies (Appendix E.4). Neither economy is re-estimated. Each changes the benchmark along aspects described below and keeps the remaining parameters at their Section 4 values.

The first, Economy HS (homogeneous subjective uncertainty), makes overconfidence and memory decay common across groups at the values estimated in Counterfactual II, so that subjective uncertainty is homogeneous. With communication, the common perceived signal variance is $k^{II}\sigma_e^2$, with σ_e^2 from Section 4. The matching private overconfidence follows from Equation (14) at the calibrated τ and σ_v^2. Removing communication raises the perceived signal variance to $k'_H\sigma_v^2$.

The second, Economy NB (no behavioral biases), shuts down overconfidence and memory decay by setting $k_g = 1$ and $\sigma_{g,\omega}^2 = 0$, so that subjective and objective uncertainty about the signal coincide. Households' signal variance, perceived and actual, equals σ_e^2 with communication and rises to σ_v^2 without it.

Because these economies differ from the benchmark, so do their optimal inflation rates with communication. These rates are optimal policies in different economies, not policy errors arising from the central bank's misperception. In particular, Economy HS's 4.604% differs from the 5.04% in Table 2, where the central bank misperceives the benchmark economy. Here in Economy HS, households actually have homogeneous biases and the central bank understands their expectation formation. What matters is the change in optimal inflation within each economy when central bank communication is removed.

Table 4 shows that central bank communication has much smaller effects in both comparison economies. In Economy HS, communication lowers optimal inflation by 7.7 basis points, less than half the benchmark effect. In the benchmark, lower-income households are less overconfident about their private information. They therefore perceive it as less precise, exhibit greater subjective uncertainty, and place more weight

Table 4. *Central bank communication and optimal inflation*

Economy	π^{opt}		Effect of communication (bp)	Share of inflation bias removed
	With communication	Without communication		
Benchmark	5.235%	5.433%	19.8	5.8%
HS: homogeneous subjective uncertainty	4.604%	4.681%	7.7	2.9%
NB: no behavioral biases	3.065%	3.074%	0.9	0.8%

Note: The effect of communication is π^{opt} without communication minus π^{opt} with communication. The share of inflation bias removed divides this effect by $\pi^{opt} - \pi^*$ without communication, with $\pi^* = 2\%$. In every economy, the central bank knows households' expectation formation. The benchmark rate with communication is the optimum in Table 2, reported to three decimals. Perceived signal variances in each economy are reported in Table A.9. All calculations use equal population shares.

on the central bank signal. The weight on the central bank signal also declines with income, from 25.7% for Group 1 to 7.1% for Group 5, shown in Table A.8. Removing communication takes away a larger share of their information and produces a large increase in their subjective uncertainty. Because these households also have the greatest initial uncertainty and overestimation, the loss of communication has a large effect on aggregate expectations. Economy HS removes this declining reliance on central bank communication: every group relies on the central bank signal to the same extent, and removing communication raises every group's uncertainty by the same moderate amount. No group experiences a particularly large increase in overestimation. Heterogeneous subjective uncertainty therefore amplifies the aggregate consequences of communication: the changes in uncertainty and overestimation are concentrated among households that initially overestimate inflation the most.

In Economy NB, optimal inflation rises by less than 1 basis point when the communication is removed. Compared to Economy HS, Economy NB further shuts down the difference between subjective and objective uncertainty. When households correctly sees its signal as very noisy, they put little weight on it and rely primarily on their priors. A reduction in signal precision therefore has little effect on their subjective uncertainty. In the benchmark, by contrast, overconfidence makes signals look far more precise than they are, and memory decay makes recalled priors less precise. Both biases raise the weight on new information, making subjective uncertainty more responsive to signal precision. This comparison shows that the effect of communication depends on the behavioral biases that determine how households use their signals.

These comparisons show that the effect of central bank communication depends on both the heterogeneity in subjective uncertainty and the difference between subjective and objective uncertainty. When subjective uncertainty is homogeneous, its effect is about 60% smaller. When households perceive uncertainty correctly, communication has little effect on the inflation rate chosen by the central bank. Recognizing behav-

ioral biases and their heterogeneity across income groups is therefore important for evaluating central bank communication.

6 Conclusion

This paper shows why distinguishing subjective uncertainty from robust preferences matters for interpreting household inflation expectations. Greater inflation overestimation among lower-income households need not imply a stronger concern about model misspecification. Using the Survey of Consumer Expectations, we document that subjective uncertainty declines with income alongside overestimation. Allowing both sources of heterogeneity to operate, our structural estimates yield an inverted-U profile of robust preferences. The declining overestimation therefore reflects the joint variation in subjective uncertainty and robust preferences.

The model connects subjective uncertainty to how households process information. Overconfidence lowers subjective uncertainty, whereas memory decay raises it; both biases generate overreaction to new information. Their opposing income gradients help explain why approximately flat overreaction coexists with declining subjective uncertainty and overestimation. Similar degrees of overreaction thus conceal substantial differences in expectation formation. Measuring subjective uncertainty helps reveal these differences and distinguish their contributions to household forecasts.

The policy analysis establishes the significance of this distinction. In our estimated economy, assuming common robust preferences shifts the chosen inflation rate from the optimum by approximately 2 basis points, while assuming common subjective uncertainty shifts it by approximately 19 basis points, with substantial losses concentrated among lower-income households. Accounting for heterogeneous subjective uncertainty is therefore particularly important for the policy prescription and its distributional consequences.

An important direction for future research is to endogenize the relationship between income and the strength of behavioral biases. We treat overconfidence and memory decay as group-specific characteristics and estimate how they covary with income, without explaining why these relationships arise. Models in which financial circumstances affect the resources available for retaining information or the incentives underlying confidence could provide foundations for these patterns. Such models could also examine how changes in income alter subjective uncertainty and expectation formation, and whether these responses modify the effects of monetary policy.

References

Afrouzi, H., S. Y. Kwon, A. Landier, Y. Ma, and D. Thesmar (2023). Overreaction in expectations: Evidence and theory. *Quarterly Journal of Economics 138*(3), 1713–1764.

Angeletos, G.-M. and J. La'O (2020). Optimal monetary policy with informational frictions. *Journal of Political Economy 128*(3), 1027–1064.

Argente, D. and M. Lee (2021). Cost of living inequality during the Great Recession. *Journal of the European Economic Association 19*(2), 913–952.

Armantier, O., G. Koşar, R. Pomerantz, D. Skandalis, K. Smith, G. Topa, and W. van der Klaauw (2021). How economic crises affect inflation beliefs: Evidence from the COVID-19 pandemic. *Journal of Economic Behavior & Organization 189*, 443–469.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review 109*(6), 2333–2367.

Azeredo da Silveira, R., Y. Sung, and M. Woodford (2024). Optimally imprecise memory and biased forecasts. *American Economic Review 114*(10), 3075–3118.

Azeredo da Silveira, R. and M. Woodford (2019). Noisy memory and over-reaction to news. *AEA Papers and Proceedings 109*, 557–561.

Ball, L. M. and S. Mazumder (2019). A phillips curve with anchored expectations and short-term unemployment. *Journal of Money, Credit and Banking 51*(1), 111–137.

Bassetti, F., R. Casarin, and M. Del Negro (2022). A Bayesian approach for inference on probabilistic surveys. Technical Report 1025, Federal Reserve Bank of New York, Staff Reports.

Belmi, P., M. A. Neale, D. Reiff, and R. Ulfe (2020). The social advantage of miscalibrated individuals: The relationship between social class and overconfidence and its implications for class-based inequality. *Journal of Personality and Social Psychology 118*(2), 254–282.

Bhandari, A., J. Borovčka, and P. Ho (2025). Survey data and subjective beliefs in business cycle models. *Review of Economic Studies 92*(3), 1375–1437.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–2782.

Broer, T. and A. N. Kohlhas (2024). Forecaster (mis-)behavior. *The Review of Economics and Statistics 106*(5), 1334–1351.

Bruine de Bruin, W., W. van der Klaauw, J. S. Downs, B. Fischhoff, G. Topa, and O. Armantier (2010). Expectations of inflation: The role of demographic variables, expectation formation, and financial literacy. *Journal of Consumer Affairs 44*(2), 381–402.

Chen, H., X. Li, G. Pei, and Q. Xin (2024). Heterogeneous overreaction in expectation formation: Evidence and theory. *Journal of Economic Theory 218*, 105839.

Chen, H. and Y. Liu (2026). Expectation and confusion: Evidence and theory. *Journal of Economic Theory 235*, 106189.

Chernozhukov, V. and H. Hong (2003). An mcmc approach to classical estimation. *Journal of Econometrics 115*(2), 293–346.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. van Rooij (2023, July). How does consumption respond to news about inflation? field evidence from a randomized control trial. *American Economic Journal: Macroeconomics 15*(3), 109–52.

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

Cui, W. and V. Sterk (2021). Quantitative easing with heterogeneous agents. *Journal of Monetary Economics 123*, 68–90.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021). Exposure to grocery prices and inflation expectations. *Journal of Political Economy 129*(5), 1615–1639.

De Fiore, F., M. J. Lombardi, and J. Schuffels (2021). Monetary policy announcements and household expectations. Technical Report 956, BIS Working Paper.

Djeutem, E. and S. Xu (2025). Return expectations across the wealth distribution. *Journal of Mathematical Economics 118*, 103132.

Engelberg, J., C. F. Manski, and J. Williams (2009). Comparing the point predictions and subjective probability distributions of professional forecasters. *Journal of Business & Economic Statistics 27*(1), 30–41.

Farmer, L. E., E. Nakamura, and J. Steinsson (2024). Learning about the long run. *Journal of Political Economy 132*(10), 3334–3377.

Hansen, L. P. and T. J. Sargent (2005). Robust estimation and control under commitment. *Journal of Economic Theory 124*(2), 258–301.

Hansen, L. P. and T. J. Sargent (2008). *Robustness*. Princeton, NJ: Princeton University Press.

Huo, Z., M. Pedroni, and G. Pei (2024). Bias and sensitivity under ambiguity. *American Economic Review 114*(12), 4091–4133.

Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the U.S. retail sector. *Quarterly Journal of Economics 134*(2), 715–783.

Jonung, L. (1981). Perceived and expected rates of inflation in Sweden. *American Economic Review 71*(5), 961–968.

Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the household level. *Journal of Monetary Economics 91*, 19–38.

Lorenzoni, G. (2010). Optimal monetary policy with uncertain fundamentals and dispersed information. *Review of Economic Studies 77*(1), 305–338.

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

Mani, A., S. Mullainathan, E. Shafir, and J. Zhao (2013). Poverty impedes cognitive function. *Science 341*(6149), 976–980.

Michelacci, C. and L. Paciello (2024). Ambiguity aversion and heterogeneity in households' beliefs. *American Economic Journal: Macroeconomics 16*(2), 95–126.

Mincer, J. and V. Zarnowitz (1969). The evaluation of economic forecasts. In J. Mincer (Ed.), *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*. New York: National Bureau of Economic Research.

Mullainathan, S. and E. Shafir (2013). *Scarcity: Why Having Too Little Means So Much*. New York: Times Books/Henry Holt and Co.

Nagel, S. and Z. Xu (2022). Asset pricing with fading memory. *The Review of Financial Studies 35*(5), 2190–2245.

Neligh, N. (2024). Rational memory with decay. *Journal of Economic Behavior & Organization 223*, 120–145.

Paciello, L. and M. Wiederholt (2014). Exogenous information, endogenous information, and optimal monetary policy. *Review of Economic Studies 81*(1), 356–388.

Pei, G. (2024). Pessimism, disagreement, and economic fluctuations. *Journal of the European Economic Association 22*(3), 1177–1227.

Souleles, N. S. (2004). Expectations, heterogeneous forecast errors, and consumption: Micro evidence from the Michigan Consumer Sentiment Surveys. *Journal of Money, Credit and Banking 36*(1), 39–72.

Svensson, L. E. O. (1997). Optimal inflation targets, "conservative" central banks, and linear inflation contracts. *American Economic Review 87*(1), 98–114.

Vellekoop, N. and M. Wiederholt (2019). Inflation expectations and choices of households. Working Paper.

Appendix

A Data and Construction of Empirical Moments

A.1 Survey Sample and Income Groups

Our household-level data come from the Survey of Consumer Expectations (SCE), and the sample covers June 2013 through February 2024. To limit the influence of extreme responses, we follow Bordalo et al. (2020) and exclude observations for which the reported inflation forecast lies more than five interquartile ranges above or below the corresponding monthly median. Applying this filter separately by month allows the cutoff to adjust to changes over time in both the level and dispersion of reported inflation expectations.

The SCE records household income using 11 categories. Some of the categories at the bottom and top of the income distribution contain relatively few observations. We therefore combine the original categories into five broader groups that provide more comparable sample sizes while retaining economically meaningful income ranges: below \$30,000; \$30,000 to \$50,000; \$50,000 to \$75,000; \$75,000 to \$100,000; and above \$100,000. Throughout the paper, we refer to these categories as Groups 1 through 5, ordered from the lowest- to the highest-income group.

A.2 Inflation-Expectation Questions and Belief Distributions

In the survey, it first asks households to provide their point forecast regarding the inflation rate over the next 12 months:

Over the next 12 months, I expect the rate of [inflation/deflation] to be ___%.

Following this question, there is another question that asks households to assign probabilities to each inflation bin:

In your view, what would you say is the percent chance that, **over the next 12 months**

- the rate of inflation will be 12% or higher ___ percent chance
- the rate of inflation will be between 8% and 12% ___ percent chance
- ...
- the rate of deflation (opposite of inflation) will be between 8% and 12% ___ percent chance
- the rate of deflation (opposite of inflation) will be 12% or higher ___ percent chance

A.3 Construction of Mincer-Zarnowitz Regressions

Table A.1. *MZ estimation results*

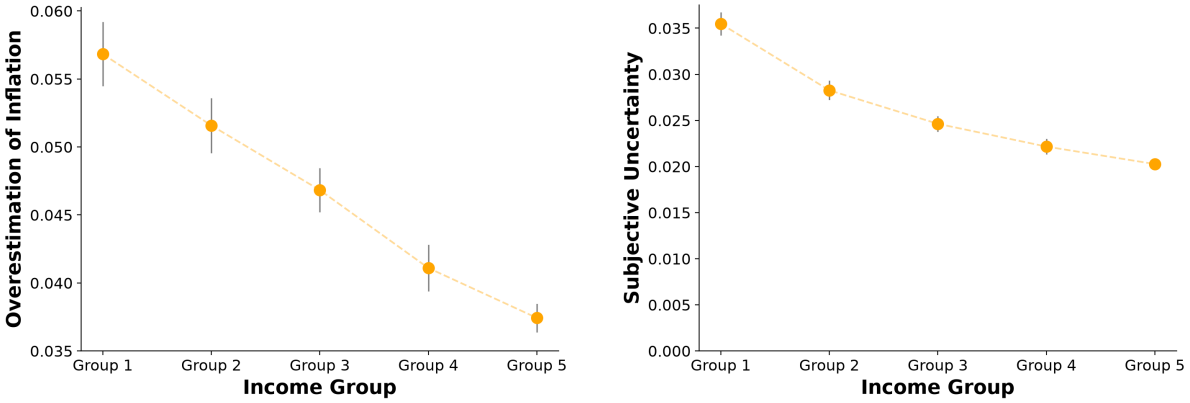
	Dependent Variable: Actual Inflation Rate				
	(1)	(2)	(3)	(4)	(5)
	Group1	Group2	Group3	Group4	Group5
Inflation Expectation	0.0498***	0.0506***	0.0594***	0.0581***	0.0732***
	(0.0087)	(0.0093)	(0.0111)	(0.0127)	(0.0133)
Constant	0.0234***	0.0244***	0.0245***	0.0261***	0.0264***
	(0.0016)	(0.0018)	(0.0018)	(0.0019)	(0.0018)
Obs.	23,789	24,789	31,715	22,533	48,130
R-sq.	0.0291	0.0210	0.0246	0.0177	0.0221

Note: This table shows results from estimating Equation (2) using the sub-sample of each income group. The sample period is from June 2013 to February 2024. Standard errors are clustered at the year-month and household levels.

B Additional Empirical Evidence

B.1 Robustness: Overestimation and Subjective Uncertainty Outside the COVID Period

In this section, we provide the empirical results using the subsample that excludes the COVID period. Following Armantier et al. (2021), we define the COVID period as 2020 Mar to 2020 Aug and exclude the data in this period. Figure A.1 shows the results. All the empirical findings are the same as our baseline results.



(a) Overestimation of inflation

(b) Subjective uncertainty

Figure A.1. *Overestimation and subjective uncertainty across income groups, excluding the COVID period. Note: Income groups are ordered from the lowest-income group (Group 1) to the highest-income group (Group 5). Panel (a) plots average overestimation of inflation, and Panel (b) plots average subjective uncertainty, using the subsample that excludes the COVID period (2020:03–2020:08, following Armantier et al. (2021)). Gray dashed lines denote 95% confidence intervals. Standard errors are clustered at the household level in both panels.*

B.2 Robustness: Subjective Uncertainty Using the EMW Method

In this section, we check whether our baseline measure of subjective uncertainty is sensitive to how the belief distribution is fitted. The baseline uses the standard deviation of the belief distribution fitted by the SCE to each household's reported bin probabilities. The distribution-mean measure used in the robustness exercises above, by contrast, assigns an ad hoc midpoint to each of the ten probability bins, including the two open-ended tail bins. Following Engelberg et al. (2009), we fit a continuous subjective distribution to each household-month response instead, using a three-tier rule based on the number of bins that receive positive probability. When three or more bins are positive, we fit a four-parameter generalized Beta distribution on a fixed truncated support of $[-20\%, 20\%]$, choosing the parameters that minimize the sum of squared differences between the model-implied cumulative distribution function and the household's reported cumulative bin probabilities at the nine interior bin boundaries, following the minimum-distance formalization in Bassetti et al. (2022).

When exactly two bins are positive, we instead fit an isosceles triangular distribution spanning their combined range, and when only one bin is positive, we use a uniform distribution over that bin's range. Each household's subjective uncertainty is then the standard deviation of its fitted distribution. Figure A.2 reports the results by income group: the declining pattern across groups is the same as in our baseline measure in Figure 1.

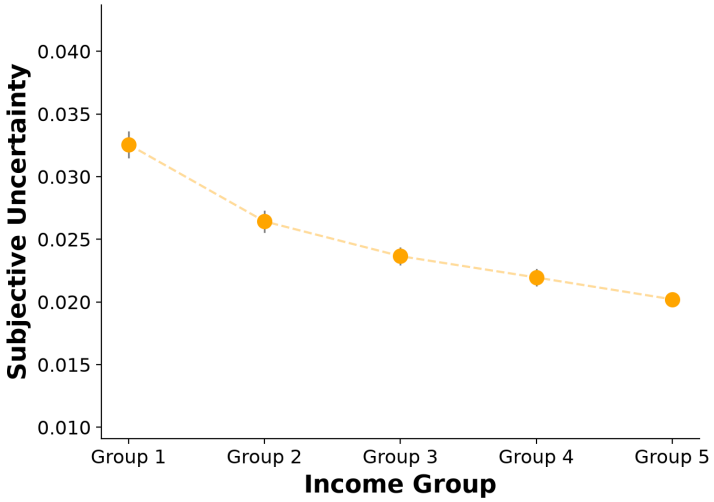


Figure A.2. *Subjective Uncertainty across Income Groups: EMW Distribution-Fitting Measure. Note: The figure reports average subjective uncertainty for each income group, measured by the standard deviation of each household's belief distribution fitted using the three-tier procedure of Engelberg et al. (2009), following the minimum-distance approach in Bassetti et al. (2022). Income groups are ordered from the lowest-income group (Group 1) to the highest-income group (Group 5). Gray dashed lines denote 95% confidence intervals. Standard errors are clustered at the household level.*

B.3 Robustness: Subjective Uncertainty Using the Interquartile Range

Our baseline measure of subjective uncertainty is the standard deviation of each household's fitted belief distribution. Because the standard deviation reflects the entire distribution, it can be sensitive to the behavior of the tails. We therefore examine whether the cross-income pattern is robust to using the interquartile range (IQR), which measures the width of the central half of the distribution. For each household, we calculate the IQR as the difference between the 75th and 25th percentiles of the fitted belief distribution and then summarize this measure separately for each income group.

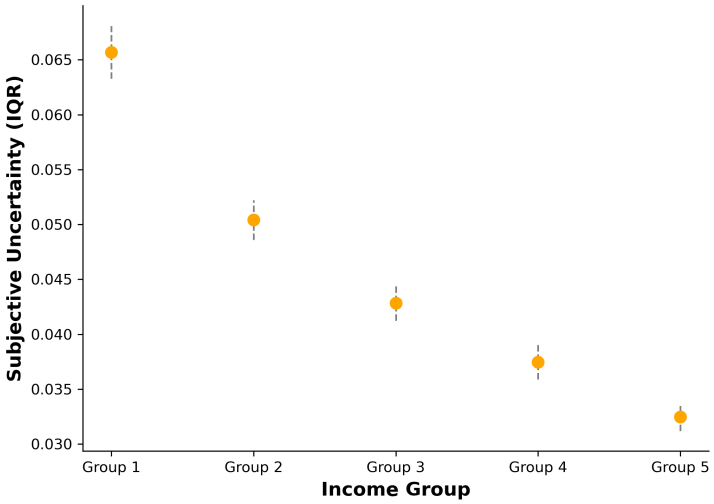


Figure A.3. *Subjective Uncertainty across Income Groups: Interquartile-Range Measure. Note: The figure reports average subjective uncertainty for each income group, measured by the difference between the 75th and 25th percentiles of each household's fitted belief distribution. Groups 1 through 5 are ordered from the lowest- to the highest-income group. The gray dashed lines indicate 95% confidence intervals. Standard errors are clustered at the household level.*

Figure A.3 shows a clear income gradient: average subjective uncertainty declines monotonically from Group 1 to Group 5. This pattern is qualitatively the same as the one obtained using the standard-deviation measure in Figure 1(b). The decline in subjective uncertainty across income groups is therefore not driven by the particular dispersion measure used in the baseline analysis or by its sensitivity to the tails of the fitted distributions.

C Calibration, Estimation, and Additional Results

C.1 Calibration of the Inflation Process

Table A.2. *Parameter Values of the Inflation Process*

Parameter		Value
Persistence of persistent component of inflation	ρ	0.85
Std dev of shocks to the persistent component	σ_ξ	0.01025
Std dev of transitory shocks	σ	0.00748
Inflation target announced by the Fed	π^*	0.02

Note: This table reports the calibrated parameters of the U.S. inflation process implied by the PCE price index.

C.2 Structural Parameter Estimates

Table A.3. *Parameter estimates*

	Income groups				
	Group 1	Group 2	Group 3	Group 4	Group 5
k	0.1877	0.0949	0.0767	0.0533	0.0520
HPDI	(0.1336, 0.2336)	(0.0768, 0.1207)	(0.0641, 0.1031)	(0.0434, 0.0745)	(0.0438, 0.0680)
σ^2_ω	0.0013	0.00078	0.00043	0.00032	0.00014
HPDI	(0.0008, 0.0023)	(0.00047,0.00115)	(0.00028, 0.00066)	(0.00020, 0.00055)	(0.00008, 0.00022)
$\alpha\chi$	24.4463	30.7419	31.9979	24.7159	19.3371
HPDI	(23.3557, 25.6564)	(29.3189, 32.2133)	(30.5442, 33.5351)	(23.1248, 26.5309)	(18.2817, 20.6848)
σ^2_ε			0.0224		
HPDI			(0.0154, 0.0319)		

Note: The estimates are obtained by sampling twenty million points in the parameter space. 68%-highest posterior density intervals (HPDI) are reported in the parentheses.

C.3 Parameter Estimates and Model Fit for the Policy Counterfactuals

Relative to the baseline, Counterfactual I targets the Mincer-Zarnowitz coefficient, subjective uncertainty, and overestimation separately for each group, while the variance of forecast errors is common. Counterfactual II targets the Mincer-Zarnowitz coefficient and subjective uncertainty at common values across income groups, while overestimation is still targeted separately for each group. Since this counterfactual assumes that households share the same behavioral biases (overconfidence and memory decay), the Mincer-Zarnowitz coefficient and subjective uncertainty are naturally common across income groups given our assumption on utility function in Section 3.

The parameter estimates are compared with the benchmark in Figure 7 in Section 5. Tables A.4 and A.5 report the targeted moments and the fit under the two restrictions.

Table A.4. *Model fit when the central bank assumes common robust preferences*

Panel A: Moments targeted as common across income groups

	Data	Model
Variance of forecast errors	0.0035	0.0034

Panel B: Moments targeted separately for each income group

	Income groups				
	Group 1	Group 2	Group 3	Group 4	Group 5
MZ coefficient: Data	0.0498	0.0506	0.0594	0.0581	0.0732
MZ coefficient: Model	0.0482	0.0470	0.0555	0.0534	0.0705
Subjective uncertainty: Data	0.0355	0.0282	0.0247	0.0224	0.0204
Subjective uncertainty: Model	0.0356	0.0296	0.0257	0.0224	0.0199
Overestimation: Data	0.0309	0.0245	0.0196	0.0124	0.0081
Overestimation: Model	0.0332	0.0229	0.0173	0.0131	0.0104

Note: This table reports the data moments used to estimate the model when the central bank assumes households share the same strength of robust preferences, and the corresponding model-implied moments. Panel A reports the moments targeted as a single common value across income groups: the variance of forecast errors. Panel B reports the Mincer-Zarnowitz coefficient, subjective uncertainty, and overestimation, each targeted separately for each income group.

Table A.5. *Model fit when the central bank assumes common subjective uncertainty*

Panel A: Moments targeted as common across income groups

	Data	Model
MZ coefficient	0.0545	0.0526
Subjective uncertainty	0.0252	0.0254
Variance of forecast errors	0.0035	0.0036

Panel B: Overestimation, targeted separately for each income group

	Income groups				
	Group 1	Group 2	Group 3	Group 4	Group 5
Data	0.0309	0.0245	0.0196	0.0124	0.0081
Model	0.0314	0.0249	0.0199	0.0125	0.0081

Note: This table reports the data moments used to estimate the model when the central bank assumes households share the same subjective uncertainty, and the corresponding model-implied moments. Panel A reports the moments targeted as a single common value across income groups: the Mincer-Zarnowitz coefficient, subjective uncertainty, and the variance of forecast errors. Panel B reports overestimation, which is targeted separately for each income group since robust preferences are allowed to vary.

Table A.6. *Group-level welfare losses when the central bank assumes common robust preferences*

	Source of Welfare Loss				
	(1)	(2)	(3)	(4)	
	Memory error	Signal error	Average FE	Inflation	Total
Group 1					
Optimal case	2.485	16.669	1.957	0.523	21.634
Counterfactual I	2.485	16.669	1.943	0.525	21.621
Contribution to additional welfare loss	0.00%	0.00%	-0.065%	0.005%	-0.059%
Group 2					
Optimal case	1.474	20.582	0.453	0.523	22.834
Counterfactual I	1.474	20.582	0.447	0.525	22.829
Contribution to additional welfare loss	0.00%	0.00%	-0.027%	0.005%	-0.022%
Group 3					
Optimal case	0.825	16.100	0.003	0.523	17.451
Counterfactual I	0.825	16.100	0.002	0.525	17.452
Contribution to additional welfare loss	0.00%	0.00%	-0.003%	0.007%	0.004%
Group 4					
Optimal case	0.562	18.792	0.246	0.523	20.123
Counterfactual I	0.562	18.792	0.251	0.525	20.129
Contribution to additional welfare loss	0.00%	0.00%	0.024%	0.006%	0.030%
Group 5					
Optimal case	0.315	11.643	1.075	0.523	13.556
Counterfactual I	0.315	11.643	1.087	0.525	13.569
Contribution to additional welfare loss	0.00%	0.00%	0.085%	0.009%	0.093%

Note: In Counterfactual I, the central bank assumes common robust preferences. The table reports group-level welfare losses measured by Equation (25) and their components. In all cases, households form expectations using the estimated group-specific overconfidence, memory decay, and robust preferences, with equal population shares. Results are obtained under the assumption that $\Psi = 1$ and $\chi = 0.1$. The "Contribution to additional welfare loss" rows report the increase in one component of welfare loss compared to the group's total loss when the central bank knows the estimates. A positive number means an increase in welfare loss when the central bank assumes common robust preferences and recognizes only the heterogeneity in overconfidence and memory decay.

C.4 Decomposition of Household Welfare Losses

Tables A.6 and A.7 report the components underlying the comparison in Table 3. Household parameters are held at their benchmark values in both cases. Memory-error and signal-error losses are therefore unchanged; the policy affects the average forecast-error loss and the direct disutility from inflation.

Table A.7. *Group-level welfare losses when the central bank assumes common subjective uncertainty*

	Source of Welfare Loss				
	(1)	(2)	(3)	(4)	
	Memory error	Signal error	Average FE	Inflation	Total
Group 1					
Optimal case	2.485	16.669	1.957	0.523	21.634
Counterfactual II	2.485	16.669	2.206	0.504	21.863
Contribution to additional welfare loss	0.00%	0.00%	1.152%	-0.092%	1.060%
Group 2					
Optimal case	1.474	20.582	0.453	0.523	22.834
Counterfactual II	1.474	20.582	0.568	0.504	22.929
Contribution to additional welfare loss	0.00%	0.00%	0.502%	-0.087%	0.415%
Group 3					
Optimal case	0.825	16.100	0.003	0.523	17.451
Counterfactual II	0.825	16.100	0.020	0.504	17.449
Contribution to additional welfare loss	0.00%	0.00%	0.098%	-0.114%	-0.016%
Group 4					
Optimal case	0.562	18.792	0.246	0.523	20.123
Counterfactual II	0.562	18.792	0.170	0.504	20.028
Contribution to additional welfare loss	0.00%	0.00%	-0.375%	-0.099%	-0.474%
Group 5					
Optimal case	0.315	11.643	1.075	0.523	13.556
Counterfactual II	0.315	11.643	0.886	0.504	13.348
Contribution to additional welfare loss	0.00%	0.00%	-1.391%	-0.147%	-1.538%

Note: In Counterfactual II, the central bank assumes common subjective uncertainty. The table reports group-level welfare losses measured by Equation (25) and their components. In all cases, households form expectations using the estimated group-specific overconfidence, memory decay, and robust preferences, with equal population shares. Results are obtained under the assumption that $\Psi = 1$ and $\chi = 0.1$. The "Contribution to additional welfare loss" rows report the increase in one component of welfare loss compared to the group's total loss when the central bank knows the estimates. A positive number means an increase in welfare loss when the central bank assumes common subjective uncertainty and recognizes only the heterogeneity in the extent of overestimation.

C.5 Model-Implied Variance of Forecast Errors

The estimation uses the variance of forecast errors in the full sample to discipline the signal-error variance, σ_e^2. We first derive the variance within an income group and then aggregate across groups. To simplify notation, we suppress group subscripts in the within-group derivation.

Iterating the forecast rule backward gives a moving-average representation in the history of private signals and memory errors. Terms that are constant within a group, such as the additive robust adjustment $\alpha\chi(\Omega + \sigma^2)$, shift the mean forecast error but not its within-group variance. The variance of the forecast error can therefore be written as

$$Var(\pi_{t+1} - \pi_{it+1|t}) =$$
$$Var\Big(\pi_{t+1}^T + \xi_{t+1} + \rho\pi_t^P - \Big[\sum_{j=0}^{\infty}(\rho A)^{j+1}\omega_{t-j} + \rho(1-A)\sum_{j=0}^{\infty}(\rho A)^j s_{t-j}\Big]\Big).$$

Substituting $s_{t-j} = \pi^P_{t-j} + \pi^T_{t-j} + e_{t-j}$ and using the autocovariance structure of the persistent inflation component yields

$$Var(\pi_{t+1} - \pi_{it+1|t}) = \frac{\sigma^2_\xi + \rho^2 A^2 \sigma^2_\omega + \rho^2(1-A)^2\sigma^2_e + [1+\rho^2(1-2A)]\sigma^2}{1-\rho^2A^2}.$$

Let V^{FE}_g denote this expression evaluated at group g's values of A_g and $\sigma^2_{\omega,g}$, and let μ^{FE}_g denote the group's mean forecast error. If r_g is group g's sample share and $\mu^{FE} = \sum_g r_g \mu^{FE}_g$ is the full-sample mean, the law of total variance gives

$$Var(\pi_{t+1} - \pi_{i,t+1|t}) = \sum_g r_g V^{FE}_g + \sum_g r_g \left(\mu^{FE}_g - \mu^{FE}\right)^2.$$

Thus, the full-sample moment combines the group-size-weighted average of within-group forecast-error variances with the between-group dispersion in mean forecast errors. This is the model-implied counterpart of the empirical moment targeted in Section 4.

C.6 Decomposing the Private Signal and Calibrating the Public Signal

Setup. In each period, household i in group g observes a private signal $x_{it} = \pi_t + v_{it}$, with $v_{it} \sim N(0, \sigma^2_v)$ i.i.d. across households and over time. It also observes a central bank signal $z_t = \pi_t + u_t$, with $u_t \sim N(0, \tau^{-1})$ i.i.d. over time and common to all households. The household perceives the variance of v_{it} as $k'_g\sigma^2_v$ and the variance of u_t correctly. Its posterior depends on the two signals only through the composite

$$s_{it} = w_g x_{it} + (1-w_g)z_t, \qquad w_g = \frac{(k'_g\sigma^2_v)^{-1}}{(k'_g\sigma^2_v)^{-1} + \tau},$$

whose perceived variance S_g is given by Equation (13). Imposing $S_g = k_g\sigma^2_e$ (Equation (14)) gives

$$k'_g = \frac{S_g}{\sigma^2_v(1-\tau S_g)}, \qquad w_g = 1 - \tau S_g,$$

which requires $\tau S_g < 1$ for every g. At the calibrated values, τS_g ranges from 0.071 (Group 5) to 0.257 (Group 1), and every $k'_g \in (0, 1]$.

The filtering recursion, the weight A_g, and subjective uncertainty depend on the signal only through S_g. So the belief updating process is consistent with the estimates in Section 4.

The composite noise has an idiosyncratic component $w_g v_{it}$ and a common component $(1-w_g)u_t$, with objective variances

$$\sigma^2_{v,g} \equiv (1-\tau S_g)^2\sigma^2_v, \qquad \sigma^2_{u,g} \equiv \frac{(1-w_g)^2}{\tau} = \tau S^2_g.$$

Variance of individual forecast errors. Both noise components are i.i.d. over time, so the derivation in Appendix C.5 applies with σ_ϵ^2 replaced by the objective variance of the composite noise, $\sigma_{v,g}^2 + \sigma_{u,g}^2$:

$$V_g^{FE} = \frac{\sigma_\xi^2 + \rho^2 A_g^2 \sigma_{\omega,g}^2 + \rho^2 (1 - A_g)^2 \left(\sigma_{v,g}^2 + \sigma_{u,g}^2\right) + \left[1 + \rho^2 (1 - 2A_g)\right] \sigma^2}{1 - \rho^2 A_g^2}. \tag{15}$$

As in Appendix C.5, the full-sample variance is

$$M \equiv \sum_g r_g V_g^{FE} + \sum_g r_g \left(\mu_g^{FE} - \mu^{FE}\right)^2. \tag{16}$$

Variance of monthly mean forecast errors. Let $\overline{FE}_{g,t}$ denote group g's mean forecast error in month t. Averaging over households removes memory errors and private noise but not communication noise:

$$\overline{FE}_{g,t} = \pi_{t+1} - \rho(1 - A_g) \sum_{j=0}^{\infty} (\rho A_g)^j \left[\pi_{t-j} + (1 - w_g) u_{t-j}\right] - \text{const}_g.$$

The same steps as in Appendix C.5 give the group-level variance

$$\bar{V}_g^{FE} = \frac{\sigma_\xi^2 + \rho^2 (1 - A_g)^2 \sigma_{u,g}^2 + \left[1 + \rho^2 (1 - 2A_g)\right] \sigma^2}{1 - \rho^2 A_g^2}. \tag{17}$$

Communication noise has mean zero, so group g's mean forecast error is μ_g^{FE}, as before. Pooling monthly mean forecast errors across groups, with group g weighted by r_g, the law of total variance gives

$$\bar{M} \equiv \sum_g r_g \bar{V}_g^{FE} + \sum_g r_g \left(\mu_g^{FE} - \mu^{FE}\right)^2. \tag{18}$$

Solution. Define $\kappa_g \equiv \rho^2 (1 - A_g)^2 / (1 - \rho^2 A_g^2)$ and $\Gamma_g \equiv \{\sigma_\xi^2 + [1 + \rho^2 (1 - 2A_g)]\sigma^2\} / (1 - \rho^2 A_g^2)$. The private-noise variance σ_v^2 does not appear in Equation (18), and τ enters it linearly through $\sigma_{u,g}^2 = \tau S_g^2$. Subtracting Equation (18) from Equation (16) removes the common terms, which include the between-group dispersion and the communication noise. The system is therefore recursive:

$$\tau = \frac{\bar{M} - \sum_g r_g \Gamma_g - \sum_g r_g \left(\mu_g^{FE} - \mu^{FE}\right)^2}{\sum_g r_g \kappa_g S_g^2}, \qquad \sigma_v^2 = \frac{M - \bar{M} - \sum_g r_g \rho^2 A_g^2 \sigma_{\omega,g}^2 / (1 - \rho^2 A_g^2)}{\sum_g r_g \kappa_g (1 - \tau S_g)^2}. \tag{19}$$

Table A.8. *Decomposition of the perceived signal variance*

	Group 1	Group 2	Group 3	Group 4	Group 5
Perceived signal variance with communication, $k_g\sigma_e^2$ ($\times 10^{-3}$)	4.203	2.124	1.718	1.193	1.163
Overconfidence about the private signal, k'_g	0.193	0.083	0.065	0.044	0.043
Perceived signal variance without communication, $k'_g\sigma_v^2$ ($\times 10^{-3}$)	5.661	2.442	1.920	1.288	1.252
Increase without communication (%)	34.7	15.0	11.8	7.9	7.7
Weight on the central bank signal (%)	25.7	13.0	10.5	7.3	7.1

Note: $\tau = 61.26$ and $\sigma_v^2 = 0.0293$ are calibrated to the full-sample variances of individual forecast errors (0.0034744) and monthly mean forecast errors (0.0003334). The perceived signal variance with communication equals its Section 4 value by construction.

With $M = 0.0034744$ and $\bar{M} = 0.0003334$, Equation (19) gives $\tau = 61.257$ and $\sigma_v^2 = 0.02934$. Table A.8 reports the implied group-specific parameters.

D Model Proofs

This section proves the paper's formal results in the order they appear in the main text. Section D.1 derives the filtering variance and the optimal forecast (Lemma 1 and Proposition 1), which every subsequent proof builds on. Section D.2 proves the observational-equivalence result (Corollary 1) as a special case of this filtering solution. Sections D.3 and D.4 use the forecast rule and filtering variance from D.1 to sign the comparative statics of overreaction (Proposition 4) and subjective uncertainty (Proposition 2), respectively; Section D.5 then combines the results of D.4 with the overestimation term from D.1 to sign its comparative statics (Proposition 3).

D.1 Proof of Lemma 1 and Proposition 1

Steady-state filtering variance and optimal forecast. The proof has two parts. The first derives the filtering recursion for $\Omega_{t|t-1}$ and shows it converges to a unique steady state Ω, proving Lemma 1. The second solves the household's max–min forecasting problem given this steady-state belief, proving Proposition 1.

Filtering and steady-state variance. $\Omega_{t|t-1}$ denotes household i's *predictive* variance about π_t^P entering period t, based on information through $t-1$; $\Omega_{t|t}$ denotes the *posterior* variance after observing m_{it} and s_{it} in period t. Because π_{t+1}^P follows the AR(1) process $\pi_{t+1}^P = \rho\pi_t^P + (1-\rho)\pi^* + \xi_{t+1}$ with $\xi_{t+1} \sim N(0, \sigma_\xi^2)$ independent of period-t information, the posterior variance $\Omega_{t|t}$ becomes next period's predictive variance according to $\Omega_{t+1|t} = \rho^2\Omega_{t|t} + \sigma_\xi^2$.

In period $t-1$, household i's belief about π_t^P follows $N(\pi_{it|t-1}^P, \Omega_{t|t-1})$. With memory decay, households' memory in period t is given by:

$$m_{it} = \pi_{it|t-1}^P + \omega_{it}, \omega_{it} \sim N(0, \sigma_\omega^2),$$

so the effective prior variance entering the period-t update is $\Omega_{t|t-1} + \sigma_\omega^2$. Combining m_{it} with the private signal s_{it} by their perceived precisions, households' posterior regarding the persistent component is given by:

$$\pi_{it|t}^P = \frac{\sigma^2 + k\sigma_e^2}{\Omega_{t|t-1} + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2} m_{it} + \frac{\Omega_{t|t-1} + \sigma_\omega^2}{\Omega_{t|t-1} + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2} s_{it},$$

$$\text{with } \Omega_{t|t} = \frac{(\Omega_{t|t-1} + \sigma_\omega^2)(\sigma^2 + k\sigma_e^2)}{\Omega_{t|t-1} + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2}.$$

Substituting $\Omega_{t+1|t} = \rho^2\Omega_{t|t} + \sigma_\xi^2$ gives the recursion in $\Omega_{t|t-1}$ alone. In the steady state, $\Omega_{t|t-1} = \Omega_{t+1|t} = \Omega$ satisfies:

$$\Omega = \rho^2 \frac{(\Omega + \sigma_\omega^2)(\sigma^2 + k\sigma_e^2)}{\Omega + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2} + \sigma_\xi^2,$$

equivalently

$$\Omega^2 + [(1-\rho^2)(\sigma^2 + k\sigma_e^2) + \sigma_\omega^2 - \sigma_\xi^2]\Omega - \sigma_\xi^2\sigma_\omega^2 - (\sigma_\xi^2 + \rho^2\sigma_\omega^2)(\sigma^2 + k\sigma_e^2) = 0. \quad (20)$$

The constant term of this quadratic is strictly negative, so it has one positive and one negative root; the positive root is the unique steady-state value, and its closed form is

$$\Omega = \frac{1}{2}\left[-B + \sqrt{B^2 + 4[\sigma_\xi^2\sigma_\omega^2 + (\sigma_\xi^2 + \rho^2\sigma_\omega^2)(\sigma^2 + k\sigma_e^2)]}\right], \qquad B \equiv (1-\rho^2)(\sigma^2 + k\sigma_e^2) + \sigma_\omega^2 - \sigma_\xi^2.$$

Household i's subjective uncertainty about π_{t+1}^P is therefore Ω, and about $\pi_{t+1} = \pi_{t+1}^P + \pi_{t+1}^T$ is $\Omega + \sigma^2$. This gives the proof of Lemma 1.

Robust distortion and optimal forecast. With the steady-state posterior belief in hand, households choose their forecast $\pi_{it+1|t}$ to solve the max–min problem in Equation (4). We first solve the inner minimization over the density distortion n_t, then the outer maximization over $\pi_{it+1|t}$, and only then impose the well-posedness restriction needed to make the resulting tilted density a valid normal distribution. The Lagrangian for the inner problem is

$$\mathcal{L} = E_t^*\left[u(\pi_{it+1|t}, \pi_{t+1})n_t + \frac{1}{\alpha}n_t \ln n_t\right] + \lambda(1 - E_t^*[n_t]),$$

with F.O.C. $u(\pi_{it+1|t}, \pi_{t+1}) + \frac{1}{\alpha}\ln n_t + \frac{1}{\alpha} - \lambda = 0$, so $n_t \propto \exp(-\alpha u(\pi_{it+1|t}, \pi_{t+1}))$; imposing $E_t^*[n_t] = 1$ pins down the constant of proportionality,

$$n_t = \frac{e^{-\alpha u(\pi_{it+1|t}, \pi_{t+1})}}{E_t^*[e^{-\alpha u(\pi_{it+1|t}, \pi_{t+1})}]}.$$

Substituting back into Equation (4), households' objective function becomes

$$\max_{\pi_{it+1|t}} -\frac{1}{\alpha}\ln E_t^*\left[\exp(-\alpha u(\pi_{it+1|t}, \pi_{t+1}))\right].$$

Using $\partial u(\pi_{it+1|t}, \pi_{t+1})/\partial \pi_{it+1|t} = \pi_{t+1} - \pi_{it+1|t}$, the first-order condition for this outer problem is

$$\frac{E_t^*\left[\exp(-\alpha u(\pi_{it+1|t}, \pi_{t+1}))(\pi_{t+1} - \pi_{it+1|t})\right]}{E_t^*\left[\exp(-\alpha u(\pi_{it+1|t}, \pi_{t+1}))\right]} = 0,$$

which rearranges to

$$\pi_{it+1|t} = E_t^*\left[\pi_{t+1}\frac{\exp(-\alpha u(\pi_{it+1|t}, \pi_{t+1}))}{E_t^*[\exp(-\alpha u(\pi_{it+1|t}, \pi_{t+1}))]}\right].$$

Households therefore forecast under a tilted posterior

$$\tilde{p}(\pi_{t+1}|m_{it}, s_{it}) \propto p^*(\pi_{t+1}|m_{it}, s_{it}) \exp(-\alpha u(\pi_{it+1|t}, \pi_{t+1})),$$

where $p^*(\cdot) = N(\mu, \Omega + \sigma^2)$ is the density under memory decay and overconfidence alone, with

$$\mu \equiv \rho \frac{\sigma^2 + k\sigma_e^2}{\Omega + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2} m_{it} + \rho \frac{\Omega + \sigma_\omega^2}{\Omega + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2} s_{it} + (1-\rho)\pi^*.$$

For $\tilde{p}$ to be a well-defined (normalizable) density, we impose $\alpha(\Omega + \sigma^2) < 1$, i.e., Assumption 1; the role of this restriction becomes clear immediately below, once we complete the square in the exponent. Using $-\alpha u(\pi_{it+1|t}, \pi_{t+1}) = \frac{\alpha}{2}(\pi_{it+1|t} - \pi_{t+1})^2 + \alpha\chi\pi_{t+1}$ and collecting terms in $\tilde{p}$ by powers of π_{t+1}:

$$\tilde{p} \sim N\left(\frac{\frac{\mu}{\Omega+\sigma^2} - \alpha\pi_{it+1|t} + \alpha\chi}{\frac{1}{\Omega+\sigma^2} - \alpha}, \frac{1}{\frac{1}{\Omega+\sigma^2} - \alpha} \right),$$

where the coefficient $\frac{1}{\Omega+\sigma^2} - \alpha$ must be strictly positive for this variance to be well defined; Assumption 1 guarantees exactly this. Solving the fixed point $\pi_{it+1|t} = \frac{\mu/(\Omega+\sigma^2) - \alpha\pi_{it+1|t} + \alpha\chi}{1/(\Omega+\sigma^2) - \alpha}$ gives

$$\pi_{it+1|t} = \rho\left(\frac{\sigma^2 + k\sigma_e^2}{\Omega + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2} m_{it} + \frac{\Omega + \sigma_\omega^2}{\Omega + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2} s_{it} \right) + (1-\rho)\pi^* + \alpha\chi(\Omega + \sigma^2),$$

which is Equation (5). This ends the proof of Proposition 1.

D.2 Proof of Corollary 1

Observational equivalence of overconfidence and memory decay. The proof uses the steady-state filtering solution derived in Section D.1. The strategy is to equate the weight placed on new information under the two restricted models, overconfidence alone ($\sigma_\omega^2 = 0$) and memory decay alone ($k = 1$), and to solve for the memory-noise variance that reproduces a given degree of overconfidence.

Both restricted models share the closed-form expression for Ω derived in the proof of Lemma 1, evaluated at the corresponding parameter restriction: write Ω^{ov} for that expression with $\sigma_\omega^2 = 0$, and Ω^{md} for the same expression with $k = 1$. Matching the weight placed on new information in the two cases, and denoting the common weight on the recalled prior by A,

$$A = \frac{\sigma^2 + k\sigma_e^2}{\Omega^{ov} + \sigma^2 + k\sigma_e^2} = \frac{\sigma^2 + \sigma_e^2}{\Omega^{md} + \sigma_\omega^2 + \sigma^2 + \sigma_e^2}.$$

Both Ω^{ov} and Ω^{md} solve Equation (20) with $\sigma^2 + k\sigma_e^2$ replaced by whichever term plays that role in the respective case: $\sigma^2 + k\sigma_e^2$ itself when $\sigma_\omega^2 = 0$, and $\sigma^2 + \sigma_e^2$ when $k = 1$. Write x for this common term, so that $A = x/(\Omega + \sigma_\omega^2 + x)$ in either case. Substituting $\Omega + \sigma_\omega^2 = x(1-A)/A$ into Equation (20) then gives $\Omega = \sigma_\xi^2 + \rho^2 x(1-A)$, for either value of x.

Overconfidence case. Here $\sigma_\omega^2 = 0$ and $x = \sigma^2 + k\sigma_e^2$, so $\Omega^{ov} = x(1-A)/A$, and this yields

$$\sigma_\xi^2 = (\sigma^2 + k\sigma_e^2)\frac{(1-A)(1-\rho^2 A)}{A}.$$

Memory-decay case. Here $x = \sigma^2 + \sigma_e^2$, so $\Omega^{md} + \sigma_\omega^{*2} = x(1-A)/A$, and this yields

$$\sigma_\omega^{*2} = (\sigma^2 + \sigma_e^2)\frac{(1-A)(1-\rho^2 A)}{A} - \sigma_\xi^2.$$

Elimination. Eliminating $(1-A)(1-\rho^2 A)/A$ between the two equations gives the condition in Corollary 1:

$$\sigma_\omega^{*2} = \frac{(1-k)\sigma_\xi^2 \sigma_e^2}{\sigma^2 + k\sigma_e^2}.$$

This ends the proof of Corollary 1.

D.3 Proof of Proposition 4

Comparative statics of overreaction. The proof proceeds in four steps: (1) derive the Mincer–Zarnowitz coefficient β as a function of the weight A; (2) note its independence from robust preferences; (3) sign its response to memory decay; and (4) sign its response to overconfidence, which uses condition (R) introduced there.

Step 1: The Mincer–Zarnowitz coefficient. Define $A \equiv \frac{\sigma^2 + k\sigma_e^2}{\Omega + \sigma_\omega^2 + \sigma^2 + k\sigma_e^2}$. Household i's forecast then can be rewritten as the weighted summation of all past information:

$$\pi_{i,t+1|t} = \rho(1-A)\sum_{j=0}^{\infty}(\rho A)^j s_{it-j} + \sum_{j=0}^{\infty}(\rho A)^{j+1}\omega_{it-j} + \text{const.}$$

Evaluating $Cov(\pi_{t+1}, \pi_{it+1|t})$ and $Var(\pi_{it+1|t})$ termwise from this representation gives the M-Z coefficient β:

$$\begin{aligned}\beta &= \frac{Cov(\pi_{t+1}, \pi_{it+1|t})}{Var(\pi_{it+1|t})} \\ &= \frac{\frac{\sigma_\xi^2}{1-\rho^2}(1-A)(1-\rho^2 A^2)}{A^2(1-\rho^2 A)\sigma_\omega^2 + (1-A)^2[(\sigma^2+\sigma_e^2)(1-\rho^2 A) + \frac{\sigma_\xi^2}{1-\rho^2}(1+\rho^2 A)]}. \end{aligned} \tag{21}$$

Step 2: Independence from robust preferences. The parameter α governing the strength of robust preferences does not appear in Equation (21): robust preferences add a constant

term to the forecast (Equation (5)), which shifts the mean of $\pi_{it+1|t}$ but leaves its covariance with π_{t+1} and its variance unchanged. Hence β is independent of α, proving part (iii).

Step 3: Effect of memory decay. We show that β is decreasing in the extent of memory decay. Throughout, for convenience, let $c \equiv \sigma^2 + k\sigma_e^2$, $d \equiv \sigma^2 + \sigma_e^2$, and $s \equiv \sigma_\xi^2/(1-\rho^2)$, so that $A = c/(\Omega + \sigma_\omega^2 + c)$ and Equation (21) becomes

$$\frac{1}{\beta} = \underbrace{\frac{\sigma_\omega^2}{s}\psi(A)}_{\equiv T_1} + \underbrace{\frac{d}{s}\frac{(1-\rho^2 A)(1-A)}{1-\rho^2 A^2}}_{\equiv T_2} + \underbrace{\frac{(1+\rho^2 A)(1-A)}{1-\rho^2 A^2}}_{\equiv T_3},$$

$$\psi(A) \equiv \frac{A^2(1-\rho^2 A)}{(1-A)(1-\rho^2 A^2)}.$$

We proceed in two steps: we first ask how T_1, T_2, T_3 vary with the weight A, and then how A itself varies with σ_ω^2 through Ω; the chain rule then combines the two.

Differentiating $1/\beta = T_1 + T_2 + T_3$ with respect to A,

$$\begin{aligned}
\frac{\partial T_1}{\partial A} &= \frac{\sigma_\omega^2}{s}\psi'(A) > 0, \\
\frac{\partial T_2}{\partial A} &= -\frac{d}{s}\cdot\frac{\rho^2(1+\rho^2)\left(A - \frac{2}{1+\rho^2}\right)^2 + \frac{(1-\rho^2)^2}{1+\rho^2}}{(1-\rho^2 A^2)^2} < 0, \\
\frac{\partial T_3}{\partial A} &= \frac{-(1-\rho^2)(1+\rho^2 A^2)}{(1-\rho^2 A^2)^2} < 0.
\end{aligned}$$

For convenience, we denote

$$B \equiv (1-\rho^2)c + \sigma_\omega^2 - \sigma_\xi^2$$

and, using Δ for the discriminant (reserving $D(A)$ below for a different, unrelated polynomial),

$$\Delta \equiv B^2 + 4\left[\sigma_\xi^2\sigma_\omega^2 + (\sigma_\xi^2 + \rho^2\sigma_\omega^2)c\right],$$

such that Equation (20) from the proof of Lemma 1, written in terms of c, becomes $\Omega^2 + B\Omega - \sigma_\xi^2\sigma_\omega^2 - (\sigma_\xi^2 + \rho^2\sigma_\omega^2)c = 0$, and Ω's expression can be written as

$$\Omega = \frac{-B + \sqrt{\Delta}}{2}, \qquad \text{i.e.,} \qquad 2\Omega + B = \sqrt{\Delta}.$$

Differentiating the quadratic implicitly gives

$$\frac{\partial \Omega}{\partial \sigma_\omega^2} = \frac{\sigma_\xi^2 + \rho^2 c - \Omega}{\sqrt{\Delta}}, \tag{22}$$

and since $A = c/(\Omega + \sigma_\omega^2 + c)$,

$$\frac{\partial A}{\partial \sigma_\omega^2} = -\frac{c}{(\Omega + \sigma_\omega^2 + c)^2}\left(\frac{\partial \Omega}{\partial \sigma_\omega^2} + 1\right) = -\frac{c}{(\Omega + \sigma_\omega^2 + c)^2} \cdot \frac{\sigma_\xi^2 + \rho^2 c + \Omega + B}{\sqrt{\Delta}} = -\frac{A}{\sqrt{\Delta}} < 0,$$

where the last equality uses $\sigma_\xi^2 + \rho^2 c + \Omega + B = \Omega + \sigma_\omega^2 + c$, which follows from the definition of B.

Since T_1 also depends on σ_ω^2 directly, and not only through A, the total derivative is

$$\frac{\partial(1/\beta)}{\partial \sigma_\omega^2} = \underbrace{\frac{1}{s}\psi(A)}_{\text{direct effect via } T_1} + \left(\frac{\partial T_1}{\partial A} + \frac{\partial T_2}{\partial A} + \frac{\partial T_3}{\partial A}\right)\frac{\partial A}{\partial \sigma_\omega^2}.$$

Multiplying by $Q \equiv c(1-A)^2(1-\rho^2 A^2)^3/A^2$, we get:

$$\begin{aligned}\frac{\partial(1/\beta)}{\partial \sigma_\omega^2} &\propto \Phi(A) \\ &\equiv c\,\rho^2\,(1-A)^2(1-\rho^2A^2)^2 + (d-c)(1-A)^2\,D(A) + (1-\rho^2)\,s\,(1-\rho^2A^2)^2,\end{aligned} \tag{23}$$

where $D(A) \equiv \rho^2(1+\rho^2)A^2 - 4\rho^2 A + (1+\rho^2)$ is exactly the numerator already shown positive above when computing $\partial T_2/\partial A$, i.e., $D(A) = \rho^2(1+\rho^2)\left(A - \frac{2}{1+\rho^2}\right)^2 + \frac{(1-\rho^2)^2}{1+\rho^2} > 0$. Since also $d \geq c$, every term of $\Phi(A)$ is nonnegative, so $\partial(1/\beta)/\partial\sigma_\omega^2 > 0$ and $\partial\beta/\partial\sigma_\omega^2 < 0$: stronger memory decay lowers β, i.e., increases overreaction. This proves part (i).

Step 4: Effect of overconfidence. We show that the Mincer-Zarnowitz coefficient increases with k; since stronger overconfidence corresponds to a lower k, this proves part (ii).

Differentiating Ω and A w.r.t. k gives

$$\frac{\partial \Omega}{\partial k} = \frac{\sigma_e^2[(\sigma_\xi^2 + \rho^2\sigma_\omega^2) - (1-\rho^2)\Omega]}{\sqrt{\Delta}}, \tag{24}$$

$$\frac{\partial A}{\partial k} = \frac{\sigma_e^2(\Omega + \sigma_\omega^2) - c\,\frac{\partial \Omega}{\partial k}}{(\Omega + \sigma_\omega^2 + c)^2} = \frac{\sigma_e^2(\sigma_\omega^2 + \sigma_\xi^2)}{\sqrt{\Delta}\,(\Omega + \sigma_\omega^2 + c)} > 0.$$

It remains to sign $\partial\beta/\partial A$. We use the condition that the recalled prior, net of recall noise, remains informative about π_t^P: its associated variance cannot exceed the unconditional variance of π_t^P itself, since conditioning on additional information should not raise variance relative to having none. Writing $s \equiv \sigma_\xi^2/(1-\rho^2)$ for this unconditional variance,

$$\Omega + \sigma_\omega^2 \leq s.$$

Since $\Omega + \sigma_\omega^2 = c(1-A)/A$, this becomes

$$c(1-A) \leq sA. \tag{R}$$

Substituting $\Omega + \sigma_\omega^2 = c(1-A)/A$ into the steady-state fixed point $\Omega = \rho^2 c(\Omega + \sigma_\omega^2)/(\Omega + \sigma_\omega^2 + c) + \sigma_\xi^2$ gives $\Omega = \sigma_\xi^2 + \rho^2 c(1-A)$, and subtracting this from $\Omega + \sigma_\omega^2 = c(1-A)/A$, using $\sigma_\xi^2 = (1-\rho^2)s$,

$$\sigma_\omega^2(A) = \frac{c(1-A)(1-\rho^2 A)}{A} - (1-\rho^2)s.$$

From Equation (21), write $\beta(A, \sigma_\omega^2) = N(A)/\mathcal{D}(A, \sigma_\omega^2)$ with

$$N(A) = s(1-A)(1-\rho^2 A^2), \qquad \mathcal{D}(A, \sigma_\omega^2) = \sigma_\omega^2 E(A) + F(A),$$

$$E(A) = A^2(1-\rho^2 A), \qquad F(A) = (1-A)^2[d(1-\rho^2 A) + s(1+\rho^2 A)].$$

By the quotient rule, $\partial\beta/\partial A$ has the sign of $N'\mathcal{D} - N\mathcal{D}' = \sigma_\omega^2(N'E - NE') + (N'F - NF')$. Direct computation gives

$$N'E - NE' = -sAW(A), \qquad N'F - NF' = sAW(A)\Psi(A),$$

where

$$W(A) \equiv 2 - (1+3\rho^2)A + 2\rho^2 A^2 - \rho^2(1-\rho^2)A^3,$$

and, with $D(A)$ as in Equation (23),

$$\Psi(A) \equiv \frac{(1-A)^2[dD(A) + s(1-\rho^2)(1+\rho^2 A^2)]}{AW(A)}.$$

$W(A)$ is positive throughout $A, \rho \in (0,1)$: at the endpoints, $W(0) = 2 > 0$ and $W(1) = (1-\rho^2)^2 > 0$, and W does not change sign in between. Hence $sAW(A) > 0$ for $A \in (0,1)$, and

$$\frac{\partial \beta}{\partial A} > 0 \quad \Longleftrightarrow \quad \sigma_\omega^2 < \Psi(A).$$

We show condition (R) implies $\sigma_\omega^2(A) < \Psi(A)$. Expanding $AW(A)[\Psi(A) - \sigma_\omega^2(A)]$, writing $d = c + (d-c)$, and using the identity $(1-A)D(A) - (1-\rho^2 A)W(A) = -(1-\rho^2)(1-\rho^2 A^2)^2$ collapses the expression to

$$AW(A)[\Psi(A) - \sigma_\omega^2(A)] = (1-\rho^2)(1-\rho^2 A^2)^2[s - c(1-A)] + (d-c)(1-A)^2 D(A).$$

The second term is nonnegative since $d \geq c$ and $D(A) > 0$. For the first, condition (R) gives $s - c(1-A) \geq s - sA = s(1-A) > 0$. Hence $\Psi(A) - \sigma_\omega^2(A) > 0$ under (R), so

$\partial\beta/\partial A > 0$.

Combining this with $\partial A/\partial k > 0$,

$$\frac{\partial \beta}{\partial k} = \frac{\partial \beta}{\partial A} \cdot \frac{\partial A}{\partial k} > 0.$$

Since stronger overconfidence corresponds to a lower k, $\partial\beta/\partial k > 0$ means that stronger overconfidence lowers β and increases overreaction. This ends the proof of part (ii), and with it the proof of Proposition 4.

D.4 Proof of Proposition 2

Comparative statics of subjective uncertainty. The proof signs $\partial\Omega/\partial k$, $\partial\Omega/\partial\sigma_\omega^2$, and $\partial\Omega/\partial\alpha$ in that order, matching parts (i)–(iii) of the proposition. Throughout, $c \equiv \sigma^2 + k\sigma_e^2$ and $B \equiv (1-\rho^2)c + \sigma_\omega^2 - \sigma_\xi^2$ are inherited from the proof of Proposition 4 in Section D.3 above.

Step 1: A root-sign device. Ω is the positive root of

$$G(\Omega) \equiv \Omega^2 + B\Omega - \sigma_\xi^2\sigma_\omega^2 - (\sigma_\xi^2 + \rho^2\sigma_\omega^2)c = 0,$$

and G's constant term is negative, so its other root is negative; consequently, for any x, $x > \Omega$ is equivalent to $G(x) > 0$ for any $x \geq 0$. This lets us sign $\partial\Omega/\partial\sigma_\omega^2$ and $\partial\Omega/\partial k$ by evaluating G at the candidate threshold appearing in Equations (22) and (24), rather than by working with Ω's closed form directly.

Step 2: Effect of overconfidence. By Equation (24),

$$\frac{\partial \Omega}{\partial k} > 0 \quad \Longleftrightarrow \quad \Omega < \frac{\sigma_\xi^2 + \rho^2\sigma_\omega^2}{1-\rho^2} \quad \Longleftrightarrow \quad G\left(\frac{\sigma_\xi^2 + \rho^2\sigma_\omega^2}{1-\rho^2}\right) > 0.$$

Evaluating G at this point gives

$$G\left(\frac{\sigma_\xi^2 + \rho^2\sigma_\omega^2}{1-\rho^2}\right) = \frac{\rho^2(\sigma_\xi^2 + \sigma_\omega^2)^2}{(1-\rho^2)^2} > 0,$$

so $\partial\Omega/\partial k > 0$. Since stronger overconfidence corresponds to a lower k, subjective uncertainty $\Omega + \sigma^2$ falls as overconfidence strengthens. This proves part (i).

Step 3: Effect of memory decay. By Equation (22),

$$\frac{\partial \Omega}{\partial \sigma_\omega^2} > 0 \quad \Longleftrightarrow \quad \sigma_\xi^2 + \rho^2 c > \Omega \quad \Longleftrightarrow \quad G(\sigma_\xi^2 + \rho^2 c) > 0.$$

Evaluating G at this point gives

$$G(\sigma_\xi^2 + \rho^2 c) = \rho^2 c^2 > 0,$$

so $\partial\Omega/\partial\sigma_\omega^2 > 0$: subjective uncertainty rises as memory decay strengthens. This proves part (ii).

Step 4: Independence from robust preferences. The fixed-point equation $G(\Omega) = 0$, equivalently Equation (20), depends only on $k, \sigma_\omega^2, \rho, \sigma^2, \sigma_e^2, \sigma_\xi^2$; the robust-preference parameter α does not appear. Hence Ω, and with it subjective uncertainty $\Omega + \sigma^2$, does not vary with α: robust preferences shift the optimal forecast (Equation (5)) without changing the variance of the subjective distribution from which that forecast is computed. This proves part (iii) and ends the proof of Proposition 2.

D.5 Proof of Proposition 3

Comparative statics of pessimistic bias. Write $b \equiv \alpha\chi(\Omega + \sigma^2)$ for the pessimistic bias, and differentiate b directly rather than repeating the full expression at each step.

For parts (i) and (ii), we hold $\alpha\chi$ fixed and vary k and σ_ω^2, so that

$$\frac{\partial b}{\partial k} = \alpha\chi\frac{\partial\Omega}{\partial k}, \qquad \frac{\partial b}{\partial\sigma_\omega^2} = \alpha\chi\frac{\partial\Omega}{\partial\sigma_\omega^2}.$$

By Proposition 2, $\partial\Omega/\partial k > 0$ and $\partial\Omega/\partial\sigma_\omega^2 > 0$; since $\alpha > 0$ and $\chi > 0$ when households have robust preferences, both derivatives above are positive. Since stronger overconfidence corresponds to a lower k, $\partial b/\partial k > 0$ means the bias falls as overconfidence strengthens, giving (i); $\partial b/\partial\sigma_\omega^2 > 0$ means the bias rises as memory decay strengthens, giving (ii).

For part (iii), we hold k, σ_ω^2, and χ fixed and vary α. By Proposition 2, part (iii), Ω is independent of α, so

$$\frac{\partial b}{\partial\alpha} = \chi(\Omega + \sigma^2) > 0,$$

since $\chi > 0$ and $\Omega + \sigma^2 > 0$: the bias rises with the strength of robust preferences. This ends the proof of Proposition 3.

E Supplementary Derivations for the Monetary-Policy Application

This section derives the optimal policy in Section 5, the central bank's excess loss under either restricted interpretation, and the decomposition of household welfare losses reported in Tables A.6 and A.7.

E.1 Derivation of Equation 11

Substituting Equation (10) into the Phillips curve gives $l = \theta[(1-\bar{B})\pi - \bar{A}\pi^* - \bar{C}]$. The central bank's first-order condition is

$$\pi - \pi^* + \Psi\theta^2(1-\bar{B})\left[(1-\bar{B})\pi - \bar{A}\pi^* - \bar{C}\right] = 0.$$

Solving for inflation yields

$$\pi^{\text{opt}} = \frac{1+\Psi\theta^2(1-\bar{B})\bar{A}}{1+\Psi\theta^2(1-\bar{B})^2}\pi^* + \frac{\Psi\theta^2(1-\bar{B})}{1+\Psi\theta^2(1-\bar{B})^2}\bar{C},$$

which is Equation (11). The second derivative of the loss function is $1+\Psi\theta^2(1-\bar{B})^2 > 0$, so this policy is the unique minimizer.

E.2 Policy and Excess Loss under a Restricted Interpretation

Let $(\bar{A}^j, \bar{B}^j, \bar{C}^j)$ denote the aggregates perceived by the central bank under Counterfactual $j \in \{I, II\}$. The policy chosen under that interpretation is

$$\pi^j = \frac{1+\Psi\theta^2(1-\bar{B}^j)\bar{A}^j}{1+\Psi\theta^2(1-\bar{B}^j)^2}\pi^* + \frac{\Psi\theta^2(1-\bar{B}^j)}{1+\Psi\theta^2(1-\bar{B}^j)^2}\bar{C}^j.$$

Households continue to form expectations according to the true heterogeneous process. The realized central bank loss is therefore

$$L_{\text{CB}}(\pi^j) = \frac{1}{2}(\pi^j - \pi^*)^2 + \frac{\Psi\theta^2}{2}\left[(1-\bar{B})\pi^j - \bar{A}\pi^* - \bar{C}\right]^2.$$

An exact quadratic expansion around π^{opt} has no linear term because the benchmark first-order condition holds. Hence

$$L_{\text{CB}}(\pi^j) - L_{\text{CB}}(\pi^{\text{opt}}) = \frac{1}{2}\left[1+\Psi\theta^2(1-\bar{B})^2\right](\pi^j - \pi^{\text{opt}})^2,$$

Thus, the central bank's excess loss is proportional to the squared distance from the optimal policy. Excess loss is nonnegative and equals zero if and only if the perceived model induces the same policy as the benchmark. The perceived aggregates need not coincide individually with the true aggregates for this equality to hold.

E.3 Decomposition of Household Welfare Losses

Household welfare loss is the negative of utility in Equation (3). Write the group-average forecast and within-group variance as

$$\pi_g^e(\pi) = \hat{A}_g \pi^* + (1 - \hat{A}_g)\pi + (\alpha\chi)_g \Xi_g,$$
$$V_g = \hat{A}_g^2 \sigma_{g,\omega}^2 + (1 - \hat{A}_g)^2 \sigma_e^2.$$

The memory and signal errors have zero mean and are mutually independent. Averaging the squared forecast error over households in group g, normalized to unit mass, therefore gives

$$\int (\pi_{gi} - \pi)^2 \, di = V_g + [\pi_g^e(\pi) - \pi]^2.$$

Substitution into the utility function yields Equation (12). Expanding V_g and $\pi_g^e(\pi)$ gives the four components reported in the detailed welfare tables:

$$\begin{aligned} L_g(\pi) = \; & \underbrace{\frac{1}{2}\hat{A}_g^2 \sigma_{g,\omega}^2}_{\text{memory error}} + \underbrace{\frac{1}{2}(1 - \hat{A}_g)^2 \sigma_e^2}_{\text{signal error}} \\ & + \underbrace{\frac{1}{2}\left[\pi - \hat{A}_g \pi^* - (1 - \hat{A}_g)\pi - (\alpha\chi)_g \Xi_g\right]^2}_{\text{average forecast error}} + \underbrace{\chi\pi}_{\text{inflation}}. \end{aligned} \tag{25}$$

The percentage change in household welfare loss reported in Table 3 is

$$100 \times \frac{L_g(\pi^j) - L_g(\pi^{\text{opt}})}{L_g(\pi^{\text{opt}})}, \qquad j \in \{I, II\}.$$

A negative value indicates a welfare gain for that group. Unlike the central bank's excess loss, this group-level change need not be nonnegative, because the central bank minimizes Equation (8) rather than each group's household loss.

E.4 Policy Environment with Central Bank Communication

Forecasts and aggregate expectations. With the decomposed signal, household i in group g combines its recalled prior m_{gi} with the composite $s_{gi} = w_g x_{gi} + (1 - w_g)z = \pi + w_g v_{gi} + (1 - w_g)u$. Its forecast is Equation (9), with s_i replaced by s_{gi} and with $\hat{A}_g$ and Ξ_g evaluated at the perceived signal variance S_g. Averaging within groups removes ω_{gi} and v_{gi} but not u:

$$\pi^e = \bar{A}\pi^* + \bar{B}\pi + \bar{C} + \bar{U}, \qquad \bar{U} \equiv u \sum_g r_g (1 - \hat{A}_g)(1 - w_g).$$

The central bank chooses π before u is realized and minimizes expected loss. $\bar{U}$ has mean zero and enters the Phillips curve additively, so

$$\mathrm{E}\left[l^2\right] = \theta^2\left[(1-\bar{B})\pi - \bar{A}\pi^* - \bar{C}\right]^2 + \theta^2\,\mathrm{Var}(\bar{U}),$$

and the last term does not depend on π. The first-order condition is the same as in Appendix E.1, so π^{opt} is given by Equation (11). Communication noise raises the central bank's expected loss by the constant $\Psi\theta^2\mathrm{Var}(\bar{U})/2$ but does not affect its policy choice. Optimal policy depends on households' information only through the perceived signal variances S_g.

Shutting down communication. As $\tau \to 0$, $w_g \to 1$, $s_{gi} \to x_{gi}$, $\mathrm{Var}(\bar{U}) \to 0$, and $S_g \to k'_g\sigma_v^2$.

Comparison economies. Table A.9 summarizes the three economies. In all three, the robustness terms $(\alpha\chi)_g$, the inflation process, the policy parameters, and the population shares are at their benchmark values, and the central bank knows households' expectation formation.

Table A.9. *Economies used in Section 5.5*

Economy	Overconfidence and memory decay	Perceived signal variance with communication	Perceived signal variance without communication
Benchmark	k_g, $\sigma^2_{\omega,g}$ from Section 4	$k_g\sigma_e^2$	$k'_g\sigma_v^2$, with k'_g from Equation (14)
NH	Common k^{II}, $\sigma_\omega^{2,II}$ from Counterfactual II	$k^{II}\sigma_e^2$, with σ_e^2 from Section 4	$k'_H\sigma_v^2$, with $k'_H = \dfrac{k^{II}\sigma_e^2}{\sigma_v^2(1-\tau k^{II}\sigma_e^2)} = 0.065$
NB	$k_g = 1$, $\sigma^2_{\omega,g} = 0$	$\sigma_e^2 = 0.0224$	$\sigma_v^2 = 0.0293$