Target-Based Communication under Credibility-Distorted Beliefs*

Daniel Díaz†

Brandeis University

This Version: November 2026

[Latest Version]

ABSTRACT: This paper examines the social appeal of implementing target-based communication when households under-react to news and their inflation expectations are shaped by how their experienced inflation endogenously determines their credibility. I incorporate this belief mechanism into a Representative Agent (RANK) and a Heterogeneous Agent (HANK) New Keynesian model to analyse its normative implications. The main findings are summarized as follows: i) every cohort under-reacts to an inflationary shock and the welfare loss falls on the cohorts under 55, while communication mostly reaches the older cohorts, who over-predict inflation through the credibility channel; ii) full target-communication is not generally optimal, since it removes the over-prediction of the older cohorts together with the part of it that was offsetting their under-reaction; and iii) communication and traditional monetary policy are complements only in the HANK model, and communicating early yields the lowest cost. This paper contributes to the discussion by showing that the implementation of target-based communication involves distributional and cost dimensions that should be considered in its design.

*I am grateful to Raphael Schoenle for his mentorship, time and dedication to this project, and to my committee members, George Hall and Jean-Paul L'Huillier, for their guidance. I thank Catherine Mann for opening up opportunities for this project, and seminar participants at Brandeis University for helpful comments. I used Claude (Anthropic) to check the spelling, grammar, and formatting of the text. All errors are my own.

†Department of Economics, Brandeis University, 415 South Street, Waltham, MA 02453, USA. Email: **ddiaze@brandeis.edu**. Website: **ddiaze3.github.io/website**.

1. Introduction

When inflation reached four-decade highs in 2022, central banks (CB, henceforth) around the globe spoke with one voice and asked households to trust that it would come back to the target rapidly, in an attempt to prevent a de-anchoring of expectations. For example, at the 2022 Jackson Hole symposium, the Federal Reserve Chair, Jerome Powell, stated that the "overarching focus right now is to bring inflation back down to our 2 percent goal" and that "we will keep at it until we are confident the job is done" (Powell 2022).[1] Another example is the European Central Bank, which made the same request in November 2022, when its President, Christine Lagarde, argued that "the more confident the public is that inflation will return to 2% in a timely manner, the smoother the adjustment process will be" (Lagarde 2022), or the Bank of England, which in the same November told households that "it's our job to make sure that inflation returns to our 2% target" and that it expected "inflation to fall sharply from the middle of next year" (Bank of England 2022). Whether repeating the target helps households is the question of this paper.

Communicating the inflation target has become a recurrent strategy of CB, with both theoretical and empirical support. In theory, Angeletos and Sastry (2021) show that communicating the inflation target anchors expectations better than communicating the policy instrument when general-equilibrium feedbacks are strong. Empirically, Coibion et al. (2022) find that a simple statement about the 2% inflation target lowers inflation expectations as much as the full policy statement, and D'Acunto et al. (2020) find that target-based communication moves household expectations more than a statement about the interest rate.

However, these results contrast with what UK households themselves report when asked about the reasons behind the change in their inflation expectations. Available data until 2026 from the Bank of England Inflation Attitudes Survey show that, on average, 26% of them cite the level of interest rates, whereas only 16% cite the inflation target, and this gap is consistent across age, income, and education, among other characteristics. This raises the policy question of this paper: what happens to the welfare of households if the central bank communicates its target intensively?

[1] He repeated this message a year later, when he reminded the public that "two percent is and will remain our inflation target" (Powell 2023).

This paper analyses this question through the lens of distorted beliefs that incorporate how credibility is shaped by inflation experiences under diagnostic expectations (DE, henceforth). These beliefs are motivated by two additional empirical regularities for UK households documented in this paper. First, the data show an under-reaction when new information is incorporated into inflation forecasts, in line with the consensus-level evidence of Bordalo et al. (2020) but with the opposite sign to the over-reaction that DE were designed to capture. Second, households that have experienced a higher inflation rate during their lifetime tend to display a lower level of credibility, in line with Salle et al. (2023), who find that remembered price surges lower trust in the CB.

To rationalize these empirical facts, I propose a mechanism that unifies experience-based expectations with endogenous credibility in how agents form the subjective distribution of inflation from which they then retrieve their point expectations. Intuitively, agents are uncertain about the anchor around which inflation gravitates, which is either the inflation they have experienced or the inflation target. This is precisely where CB communication is introduced, since it intentionally pulls the experienced anchor toward the inflation target. Through Bayesian updating, I define the subjective probability that the agent assigns to the uncertain anchor being the inflation target as subjective credibility, which is then endogenously incorporated into the agent's predictive distribution of inflation. Finally, the agent reports the diagnostic version of this distribution, which incorporates only a fraction of the news and captures the under-reaction observed in the data.

In the model, distorted beliefs determine different ex-ante real interest rates across age cohorts, which, for the same inflationary shock, cause different consumption responses between cohorts. On impact every cohort under-reacts to the shock, so it expects too little inflation, perceives an ex-ante real rate that is too high and non-optimally postpones consumption relative to the path it would choose with correct beliefs. For the older cohorts, with a higher inflation experience and lower credibility, the forecast leans on lived inflation, so after impact they expect higher inflation, perceive a lower ex-ante real rate and non-optimally bring consumption forward. The non-optimal consumption responses relative to the path with correct beliefs, the consumption wedges henceforth, derived from the different perceived rates naturally constitute the measure with which I analyse the welfare implications of distorted beliefs.

Next, I embed the distorted beliefs mechanism into a general equilibrium

framework such as a Representative Agent New Keynesian (RANK) model and a Heterogeneous Agent New Keynesian (HANK) model, and study the welfare effects of the intensity with which the central bank communicates its target. I show that the consumption wedges are the sum of two forecast errors. First, the *diagnostic error*, which captures the households' under-reaction, is negative on impact for every cohort and beyond the reach of communication. Second, the *credibility error*, which follows from the households' inclination towards lived inflation, is positive for the older cohorts whose inflation experience diverges consistently from the target, and is the only part the CB can move.

The findings are summarized as follows. First, in both models the welfare loss from distorted beliefs falls on the young. Every cohort under-reacts to the shock on impact, but the cohorts under 55, whose inflation history shows little persistence, barely revise their forecast upwards. On the other hand, the older cohorts lose less because the persistent inflation they lived through makes their forecast follow the shock more closely via the erosion of their credibility that makes them over-predict inflation. Importantly, this over-prediction is the only part of the consumption wedges that the CB can reach with communication. The HANK model preserves this ranking with smaller wedges, because the households close to the borrowing constraint do not substitute consumption intertemporally.

The second result is concerned with the superiority of target communication discussed in Angeletos and Sastry (2021) and Coibion et al. (2022). Communicating the target removes the credibility error of every cohort, but it is not unequivocally superior. Since the diagnostic error is beyond the reach of the CB, both the benefit and the cost of communication are defined on the credibility error. The benefit is removing the over-prediction of the old cohorts. The cost is that, on impact, the credibility error was offsetting part of the under-reaction. The optimal intensity balances these two effects, and in the RANK model the cost keeps it low, at about one third.

In the HANK model this result changes because of the nature of wealth heterogeneity. After an inflationary shock, the high-MPC households close to the borrowing constraint cannot smooth the fall in labour income, so the contraction in demand is stronger and more persistent than in the RANK model. Therefore, inflation returns to the target faster, but the older cohorts keep over-estimating it for longer given their high inflation experience. Their credibility error no longer offsets their under-reaction but adds to their over-prediction, and the optimal

intensity roughly doubles.

The third result connects target communication with traditional monetary policy by analysing their interplay. For the RANK model, a more aggressive interest rate response lowers inflation and with it the diagnostic and the credibility error in the same proportion, whereas intensive target communication removes the credibility error together with the part of it that was offsetting the under-reaction. Hence, communication does not significantly increase welfare. For the HANK model the result is different, because the aggressive response brings inflation down faster while the older cohorts keep over-estimating it, so their credibility error is added to their diagnostic error. Removing it with intense target communication is worth more the more aggressive the monetary rule is, so the two policies are complements. Consistently, the lowest loss is reached with full communication and a moderately aggressive rule. Finally, the timing of communication matters. In the RANK model a short delay is slightly better than communicating immediately, since it keeps the offset on impact. In the HANK model the over-prediction right after impact is stronger, so prompt communication is best, and starting a year late is worse than not communicating at all.

This paper relates to four strands of the literature. The first studies CB credibility and makes it endogenous, dependent on past errors.[2] For instance, Hommes and Lustenhouwer (2019) model it as the fraction of the population that trusts the target rather than extrapolating past inflation, consistent with McMahon and Rholes (2025), who show experimentally that people judge the central bank by its past forecast errors. In parallel, bringing in a long-run experiential component, Malmendier and Nagel (2016) show that households form inflation expectations from the inflation they have personally lived through, and Salle et al. (2023) find that remembered price surges lower trust in the CB. This paper connects both strands by modelling credibility as a belief that each cohort updates against its own inflation history. Credibility is a probability inside the agent's mind, with the same functional form as the population fraction of Hommes and Lustenhouwer (2019), which delivers a predictive probability distribution with endogenous credibility.

The second strand studies how inflation expectations respond to news, par-

[2] In his seminal article, Blinder (2000) settles on the definition that a CB is credible if people believe it will do what it says. A reduced-form approach measures credibility from the distance between expectations and the CB inflation target, as in Bordo and Siklos (2015), from the weight of the target in expectations, as in Bomfim and Rudebusch (2000), or by asking people directly, as in Haldane and McMahon (2018).

ticularly, whether forecasters over- or under-weight it in their predictions. In an influential paper, Coibion and Gorodnichenko (2015) document that consensus forecasts under-react to news, and Bordalo et al. (2020) have shown that under-reaction at the consensus level may coexist with over-reaction at the individual level. Using the diagnostic expectations (DE) of Bordalo et al. (2018), it is possible to reconcile these over- and under-reaction differences. L'Huillier et al. (2024) and Bianchi et al. (2024) have incorporated DE into the New Keynesian model, which makes it possible to analyse this sort of distortion in general equilibrium. This paper borrows the DE operator, applies it to the predictive distribution of inflation with endogenous credibility, and uses it with a negative diagnostic parameter, so that it delivers the under-reaction of the cohort consensus observed in the data.

The third research area to which this paper connects is CB communication with the public. There is ample evidence that communication works well to anchor the expectations of the most informed agents in the economy, such as professional forecasters and financial markets. For example, Romer and Romer (2000) show that private forecasters update taking as reference the new information the Fed reveals through its policy actions, and Nakamura and Steinsson (2018) show that policy announcements move professional forecasts and asset prices at high frequency. However, communicating with the general public has been a major challenge for the CB and has often been unsuccessful, as analysed by Binder (2017) and Lamla and Vinogradov (2019), who find that policy announcements are largely unnoticed by households, and by Kumar et al. (2015), who show that even after decades of inflation targeting, firms' inflation expectations in New Zealand remain essentially unanchored. This paper takes that evidence and asks when, and for whom, communicating the target intensively improves household welfare.

Finally, this paper is also related to the literature on heterogeneous beliefs and redistribution in New Keynesian models. Auclert (2019) shows that the covariance between MPCs and URE determines the winners and losers of monetary policy. Guerreiro (2023) and Shabalina and Tzaawa-Krenzler (2024) independently study how imperfect information interacts with wealth heterogeneity by adding an additional covariance, between MPCs and beliefs, which come from disagreement or attention, respectively. This paper borrows the cohort belief structure of Guerreiro (2023) and the sequence-space representation of Auclert et al. (2021) to analyse who gains by and how much from CB communication when expectations are distorted in the way I propose.

The rest of this paper is organized as follows. Section 2 presents the stylized facts, Section 3 develops the inflation expectation formation process, Section 4 embeds it into the RANK and HANK models, Section 5 presents the results on target communication and welfare, and Section 6 concludes.

2. Empirical Motivation

This section presents the stylized facts about inflation expectations for the UK that motivate this paper. Three data sources are used. The Bank of England Inflation Attitudes Survey is the source of inflation expectations and of the measure of central bank credibility. The cohort inflation histories are built from the Consumer Price Index published by the ONS, and the wealth measure comes from the ONS Wealth and Assets Survey. Appendix D details the construction of every variable and Appendix D.4 reports descriptive statistics.

2.1. Stylized Facts

Fact 1. Under-Reaction to Inflation News.

When UK households revise their inflation expectations, realised inflation moves further than the revision.

The evidence comes from applying to the data the forecast-error-on-revision test of Coibion and Gorodnichenko (2015), which evaluates whether expectation revisions predict the subsequent forecast errors. If they do, the forecasts have not fully incorporated the news, which is evidence of a degree of information rigidity. For US professional forecasters in the period 1969–2014, they find a coefficient around 1.2 for the inflation rate, and the same positive sign for US consumers and firms, and for professional forecasters in twelve advanced economies, the UK among them.

In this spirit, Bordalo et al. (2020) apply an analogous test to an ample set of US series from the SPF and Blue Chip surveys over 1968–2016 and find that over- or under-reaction depends on the level of aggregation. On the one hand, at the aggregate level (consensus) there is under-reaction[3] as in Coibion and Gorodnichenko (2015), but at the individual level, forecasters over-react to their

[3]Similar under-reaction to macroeconomic news, including inflation, is documented for Italian and German firms by Born et al. (2025).

own revisions. They reconcile both facts with diagnostic expectations (Bordalo et al. 2018) under dispersed information.[4] Since the survey used here only allows the consensus version of the test, the relevant benchmark is the under-reaction of the consensus.

I implement the test of Equation (1) using the BoE Inflation Attitudes Survey, with quarterly data for 2010Q1–2024Q4. Since the survey is a repeated cross-section, only the consensus version of the test can be performed:

$$\underbrace{\pi_{t,t+3} - \mathbb{E}_{g,t}[\pi_{t,t+3}]}_{\mathrm{FE}_{g,t}} = \beta_{0,g} + \beta_{1,g} \underbrace{\left(\mathbb{E}_{g,t}[\pi_{t,t+3}] - \mathbb{E}_{g,t-4}[\pi_{t,t+3}]\right)}_{\mathrm{Rev}_{g,t}} + e_{g,t}. \quad (1)$$

Here $\pi_{t,t+3}$ is the realised CPI inflation over quarters t to $t+3$, measured as the average of the four quarterly annualised rates, $\mathbb{E}_{g,t}[\pi_{t,t+3}]$ is the cohort's consensus expectation over the same quarters, and $\mathbb{E}_{g,t-4}[\pi_{t,t+3}]$ is the expectation the same cohort held one year earlier for the same twelve months.[5]

In the spirit of Bordalo et al. (2022), revisions and errors are calculated on a fixed-event basis to keep a consistent comparison window. The expectation revisions, $\mathrm{Rev}_{g,t}$, show how expectations were updated today compared with one year ago for the same target window. On the other hand, forecast errors, $\mathrm{FE}_{g,t}$, show how close expectations were to realised inflation over the same window. Under-reaction is found when news that raise revisions are accompanied by positive forecast errors, which ultimately implies that the forecast fell short of fully incorporating the news.

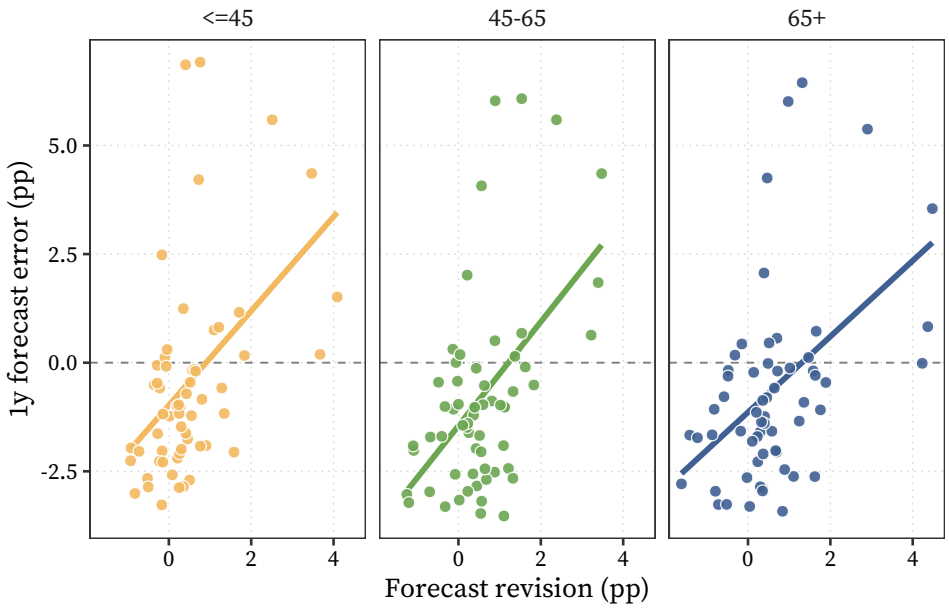
Figure 1 shows a consistent positive slope regardless of age, between 0.9 and 1.2, meaning the three cohorts under-react by a similar amount, close to the estimate of Coibion and Gorodnichenko (2015). Thus, when a cohort revises its inflation expectation upward, realised inflation turns out to be above the revised forecast. Although the revision is made in the right direction, its magnitude is lower than what a full-information response would have implied, providing evidence that only a fraction of the news was incorporated in the revised forecast. This under-

[4] They estimate a diagnosticity parameter of about 0.5 (between 0.42 and 0.59 across their three methods). When incorporating Diagnostic Expectations (DE) into the New Keynesian model, L'Huillier et al. (2024) estimate 0.73 (0.43–0.65 in robustness checks) and Bianchi et al. (2024) estimate 1.97.

[5] For the two horizons I use Q2a, which asks "How much would you expect prices in the shops generally to change over the next 12 months?", and Q2b, which follows with "And how about prices the year after next? How much would you expect prices in the shops generally to change between [month, year +2] and [month, year +3]?".

reaction is what the diagnostic parameter of Section 3 captures, with a negative sign.

FIGURE 1. **Forecast Errors and Revisions**



Note: Sample: BoE Inflation Attitudes Survey, quarterly waves 2010Q1–2024Q4, 60 waves per cohort. Slopes $\hat{\beta}_{1,g}$ of Equation (1) (WLS, robust t in parentheses): ≤45, 1.09 (3.39); 45–65, 1.20 (4.07); 65+, 0.87 (3.69).

Fact 2. Lived Inflation, Credibility and Inflation Expectations.

Cohorts that lived through higher inflation are less confident that inflation will return to the target, and cohorts that are less confident expect higher inflation.

This fact documents that cohorts that have experienced higher inflation report lower credibility and also expect a higher inflation rate. To provide evidence on this fact, the credibility measure I use is question Q18 of the BoE Inflation Attitudes Survey with sample 2011Q1–2025Q3, when Q18 is asked. This question asks "How confident are you that inflation will be close to the 2% target set by the government in two to three years' time?"[6] I define the lack of credibility of cohort g in wave t

[6]By close it means between 1% and 3%. The potential answers are "very confident", defined in the questionnaire as a better than 50% chance, "fairly confident", around 50%, "not very confident", around 25%, and "not at all confident", less than 25%.

as:

$$1 - \hat{\chi}_{g,t} \equiv \frac{\sum_{j \in \mathcal{V}_{g,t}} w_j \mathbf{1}\{\text{Q18}_j \in \{\text{not very confident, not at all confident}\}\}}{\sum_{j \in \mathcal{V}_{g,t}} w_j}, \quad (2)$$

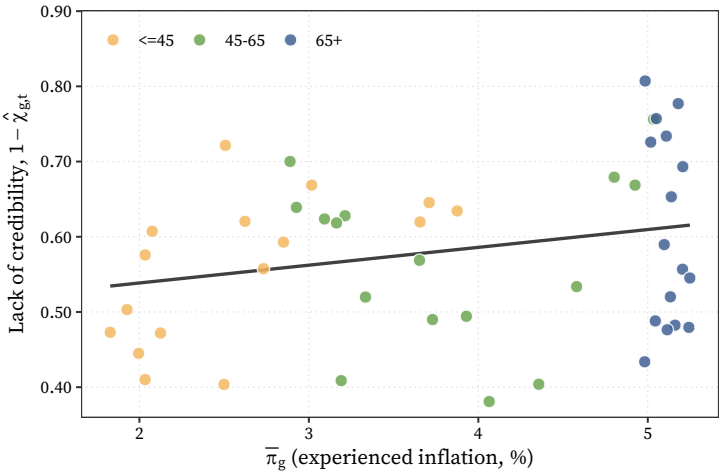
and, similarly, $\hat{\chi}_{g,t}$ is the share of the cohort that answers "very confident" or "fairly confident".[7] Hence, $\hat{\chi}_{g,t}$ is the credibility share of cohort g. It is closely related to the credibility measure of Bomfim and Rudebusch (2000), in the sense of "the fraction of the population that believes the target will be achieved", and similarly to the fraction of fundamentalists of Hommes and Lustenhouwer (2019), who believe that inflation will be at the target.

Figure 2 plots the share of each cohort that reports lack of credibility, $1 - \hat{\chi}_{g,t}$, against the experienced inflation $\bar{\pi}_g$ and it shows a positive relationship. This is consistent with the track-record view of Blinder (2000), in which credibility is earned by a history of low inflation, and with Bomfim and Rudebusch (2000), where the weight on the target falls with the past deviation of inflation from it and with Salle et al. (2023), who find that households who remember more price surges report lower trust in the central bank.

Figure 3 closes the reasoning by plotting the cohort's inflation expectations versus the self-reported credibility. At the two horizons, the slope is positive, meaning that a cohort with a higher proportion of people distrusting that inflation will return to the target expects more inflation. This is consistent with Christelis et al. (2020), who find that households with more trust in the central bank report lower and less dispersed inflation expectations, and with Haldane and McMahon (2018), who show, using the same survey, that respondents more satisfied with the Bank of England's performance hold expectations closer to the inflation target. In short, the two figures say that lived inflation lowers credibility which, in turn, raises inflation expectations. This reasoning motivates the proposed inflation expectations mechanism of Section 3.

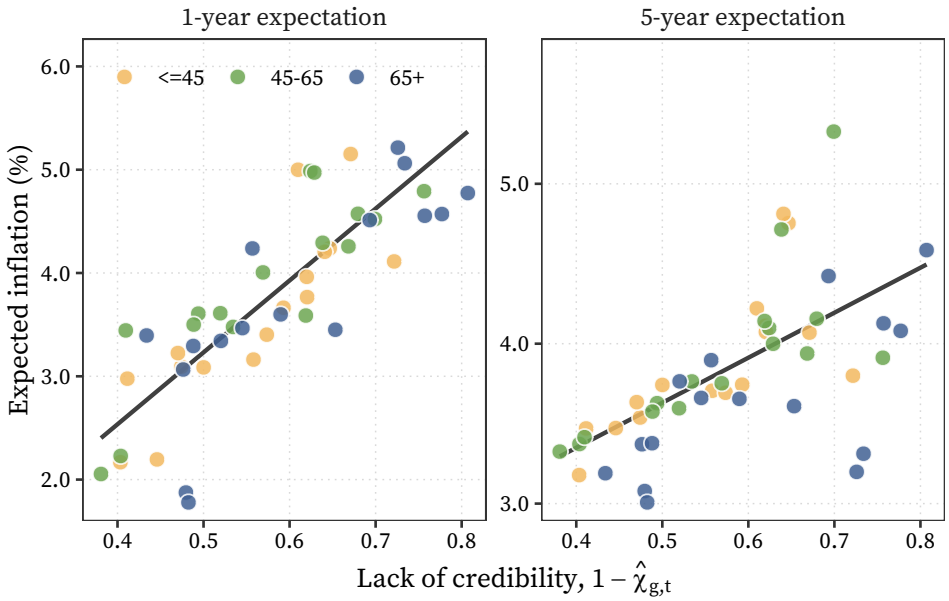
[7] The two shares are complements among valid answers, so either one defines the same $\hat{\chi}_{g,t}$.

FIGURE 2. **Credibility and Experienced Inflation**



Note: Sample: BoE Inflation Attitudes Survey, waves with Q18, 2011Q1–2025Q3, 16 waves. "don't know" excluded. OLS estimated slope of $1-\hat{\chi}_{g,t}$ is 0.0237 (4.08).

FIGURE 3. **Inflation Expectations and Credibility**

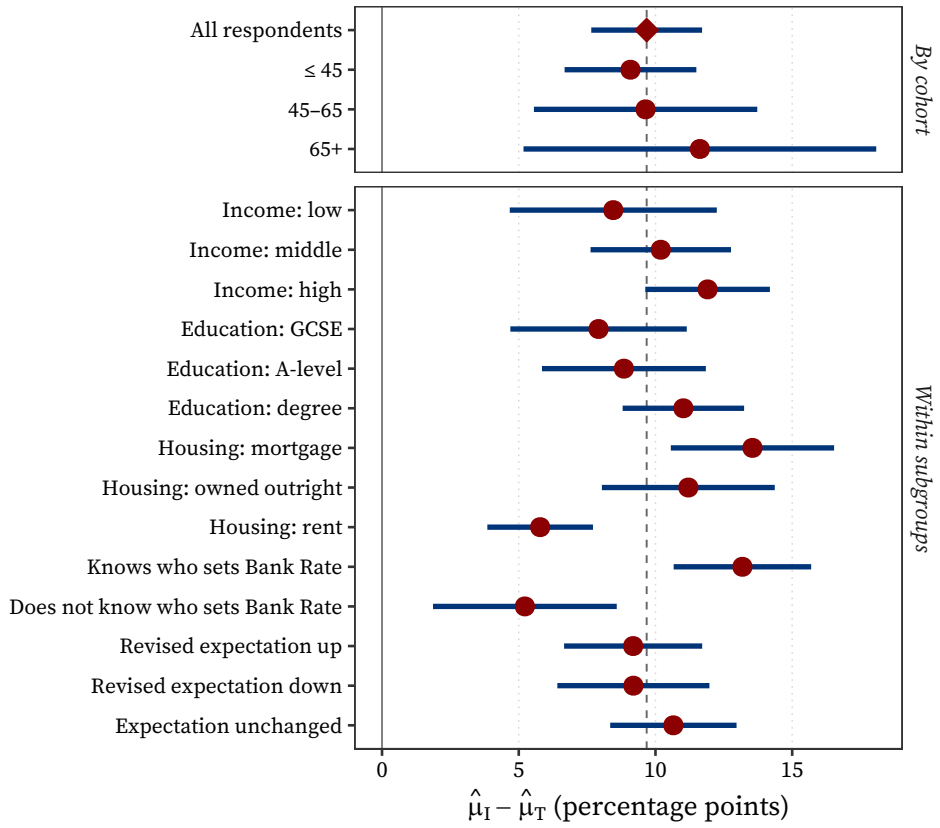


Note: Sample: BoE Inflation Attitudes Survey, waves with Q18, 2011Q1–2025Q3, 16 waves. "don't know" excluded. Slope coefficients: i) one-year horizon 6.96 (5.91); ii) five-year horizon 2.82 (4.49).

Fact 3. Instrument-Over-Target Salience.

More households report forming their inflation forecast on the interest rate than on the inflation target.

FIGURE 4. **Instrument Salience: Interest Rate versus Inflation Target**



Note: Sample: BoE Inflation Attitudes Survey, waves with Q2aiv (2017Q1, 2023Q1, 2024Q1, 2025Q1, 2025Q3, 2026Q1). Income rows use the post-2022 sample. Dots are the estimates $\hat{\mu}_I - \hat{\mu}_T$, the vertical line marks the null $\mu_I - \mu_T = 0$, and horizontal lines are the 95% confidence intervals with clustered standard errors.

To state this fact, I use question Q2aiv in the survey, which asks households who changed their expectation "What were the most important factors that led you to your expectation of prices in the shops over the next 12 months?". Across the potential reasons, I focus on the cues the central bank controls through its communication: the inflation target, $\hat{\mu}_T$, and the policy instrument, $\hat{\mu}_I$. For each case, I report the mean across cohorts in Figure 4, which shows that the weighted share citing the instrument is statistically higher than the weighted share citing

the target in every cohort and within subgroups of the population. This empirical result contrasts with what Angeletos and Sastry (2021) theoretically recommend. In their framework, communication about the target anchors expectations more effectively than communication about the instrument when general-equilibrium feedbacks are strong. However, UK households report that they intensively form their inflation expectations around the policy rate.

Two features of Figure 4 matter for the model presented in Section 3. First, the gap is positive and significant for every group, so a higher proportion of households report the interest rate as the cue they consider when forming their expectations. Second, the gap is indistinguishable across cohorts, so the salience of the instrument over the target is a common feature of the whole population rather than a characteristic of the cohorts that lived through high inflation. This is why the communication intensity of Section 3 is common across cohorts. Within subgroups, the gap is larger for households with a mortgage and for those who know who sets Bank Rate, the groups most exposed to, or informed about, the policy rate.[8]

This pattern is consistent with the evidence that central bank communication about its objectives is mainly ignored by the public. Binder (2017) and Lamla and Vinogradov (2019) find that households are mostly unaware of policy announcements, to the point that Haldane and McMahon (2018) describe the general public as the blind spot of central bank communication. Despite this unawareness of the public, there is evidence that the central bank can influence expectations with active target communication. Ehrmann et al. (2025) document that the announcement of the ECB's revised inflation target went largely unnoticed, whereas the same content delivered directly through information treatments raised the credibility of price stability. This justifies why the model of Section 3 considers the target as a feature the central bank can choose to actively communicate through the intensity ν.

Taken together, the three facts describe characteristic aspects of how UK households form inflation expectations. They under-react to news (Fact 1); the credibility they have is shaped by the inflation they have lived through, and inflation expectations are affected by it (Fact 2); and the inflation target is cited less than the interest rate as a cue when forming their expectations (Fact 3).[9] Section 3 takes

[8]The mortgage division is consistent with the unhedged interest rate exposure of Auclert (2019). The models of Section 4 does not consider this channel.

[9]Two additional facts, documented in Appendix E, support the mechanism. First, households

these facts and formalizes a theory in which agents are uncertain about the anchor around which the inflation rate gravitates. Credibility is microfounded within the agent's mind, and under-reaction is an essential component of how agents incorporate news. The result is a distorted beliefs operator that is incorporated in the upcoming sections.

3. Credibility and Diagnostic Expectations

This section formally introduces credibility-distorted beliefs, that is, how the endogenous-credibility approach of Hommes and Lustenhouwer (2019) connects with the diagnostic-expectations (DE) framework of Bordalo et al. (2020) and L'Huillier et al. (2024). The connection is not direct, since DE can be used only when a well-defined subjective distribution is available. Particularly, DE are defined as a distortion of an agent's conditional distribution relative to that same agent's lagged conditional distribution. Therefore, the diagnostic operator needs a forecaster whose beliefs at t and $t-1$ are linked by an updating rule. On the other hand, the endogenous credibility of Hommes and Lustenhouwer (2019) only provides the aggregate proportion of the population that is fundamentalist, their measure of credibility. However, it leaves unanswered how agents incorporate credibility into their subjective distribution of future inflation at the individual level. Therefore, this section proposes an approach delivering a well-defined distribution for future inflation on which the diagnostic operator can act.

Suppose cohorts indexed by g, each with a representative agent within, as Assumption 1 states.

ASSUMPTION 1 (One forecaster per cohort). *Each cohort g is a single representative forecaster. Within a cohort there is no heterogeneity of beliefs and no averaging across people. Every weight that appears below is a probability held by this one agent, not a share of a population.*

This agent believes that the inflation rate is an AR(1) process around an anchor, m, whose true value is unknown. The perceived law of future inflation is summarized by Assumption 2:

that expect more inflation report bringing forward purchases and saving more (Fact A1). Second, cohorts that lived through higher inflation perceive higher inflation today and have become less confident that higher interest rates slow down price rises (Fact A2).

ASSUMPTION 2 (Perceived Law of Future Inflation). *Cohort g believes inflation follows a first-order autoregression around an anchor m,*

$$\pi_{t+1} = \rho_g \pi_t + (1 - \rho_g)\, m + \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim N(0, \sigma_g^2) \text{ i.i.d.} \tag{3}$$

It knows ρ_g and σ_g, but it does not know m, the level inflation returns to.

Each agent balances two hypotheses about the anchor. First, under hypothesis T, the anchor is the central bank's inflation target. Second, under hypothesis E, the anchor is a communication-weighted sum of the inflation target and the lived inflation, $\bar{\pi}_g$. Introducing $\bar{\pi}_g$ as a potential anchor has ample empirical support. For instance, Malmendier and Nagel (2016) show that inflation expectations track the inflation they have lived through, as if each individual has estimated an AR(1) on its own history, assigning higher weight to recent observations. Under the latter, I assume the central bank can reduce the influence of lived inflation on the anchor and displace it toward the inflation target. Assumption 3 formally states the two hypotheses.

ASSUMPTION 3 (Hypotheses About the Anchor). *Cohort g considers two values for the anchor,*

$$m_T = \pi^*, \qquad m_{E,g} = \nu \pi^* + (1 - \nu)\, \bar{\pi}_g, \qquad \nu \in [0, 1]. \tag{4}$$

Hypothesis T says the anchor is the target. Hypothesis E says the anchor is what the cohort lived through, moved toward the target by a fraction ν that measures the intensity of the central bank's target-communication. Finally, the lived inflation is given by $\bar{\pi}_g$.

To introduce central-bank communication into the formation of expectations, I focus on the content a central bank transmits when it talks about its inflation target. As discussed in Section 2, passive announcements go largely unnoticed by households, whereas the same content delivered directly raises the perceived credibility of price stability (Ehrmann et al. 2025). This paper does not explain the inattention of the public to target communication. Instead, it assumes that the central bank can reach households if it intentionally decides to communicate its target intensively.

Accordingly, I assume the central bank directly chooses the *intensity* of its target-communication through $\nu \in [0, 1]$, where $\nu = 0$ is the baseline case in which the central bank is silent about the target and $\nu = 1$ is a target-communication

strategy at its maximal intensity. Two properties of ν are worth mentioning. First, even when ν is a common communication strategy, its effects are cohort-specific, since it acts on each cohort's own experienced gap. Second, the intensity ν does not change the target or the lived inflation, but it changes what the cohort would believe under hypothesis E. Intuitively, a household that is consistently told what the inflation target is incorporates it into its alternative hypothesis, so its experiential anchor is pulled toward π^* by the fraction ν. Thus, under maximal target communication the two hypotheses coincide at the target, and credibility becomes irrelevant for the forecast.

Households do not know which hypothesis is right and learn about it from the forecast errors of each one of them:

$$e_{h,t} \;=\; \pi_t - \rho_g\,\pi_{t-1} - (1-\rho_g)\,m_h. \tag{5}$$

Individuals incorporate hypothesis errors into future inflation beliefs via Bayesian updating. Hence, individuals calculate the log likelihood of each observation under each hypothesis given those errors as

$$\log f(\pi_t \mid \pi_{t-1}, h) \;=\; -\tfrac{1}{2}\log\!\left(2\pi\sigma_g^2\right) - \frac{e_{h,t}^2}{2\sigma_g^2}. \tag{6}$$

The learning occurs via discounting old evidence at a geometric rate ρ_U. This is the "forgetting" factor of Koop and Korobilis (2012) and essentially measures how much past performance is weighted relative to recent performance when evaluating hypothesis h. Assumption 4 formalises this idea.

ASSUMPTION 4 (Discounted Learning). *Cohort g starts with prior probability* 1/2 *on each hypothesis and updates by Bayes' rule with the discounted likelihood of each hypothesis $h \in \{T, E\}$:*

$$\ell_{h,t} \;=\; \sum_{j\geq 0} \rho_U^{\,j}\,\log f(\pi_{t-j} \mid \pi_{t-j-1}, h), \qquad \rho_U \in (0,1). \tag{7}$$

Now, consider hypothesis T (the anchor is the inflation target) and calculate its posterior using Bayes' rule, and assume that ex-ante, before any error is accumulated, the prior the agent assigns to each hypothesis is 0.5, so there is no ex-ante

bias across hypotheses. [10] I define the posterior of hypothesis T as the *subjective credibility*, which is summarized by Proposition 1.

PROPOSITION 1 (Subjective Credibility). *Under Assumptions 1–4, the probability that cohort g attaches to hypothesis T is*

$$\chi_{g,t} = \frac{e^{bU_{T,g,t-1}}}{e^{bU_{T,g,t-1}} + e^{bU_{E,g,t-1}}} = \left[1 + e^{-b\left(U_{T,g,t-1} - U_{E,g,t-1}\right)}\right]^{-1}, \qquad b = \frac{1}{2\sigma_g^2\left(1-\rho_U\right)}. \tag{8}$$

where the performance $U_{h,t}$ of hypothesis $h \in \{T, E\}$ is

$$U_{h,t} = -(1-\rho_U)\, e_{h,t}^2 + \rho_U\, U_{h,t-1}. \tag{9}$$

PROOF. See Appendix A, subsection A.1. □

There are two immediate consequences of Proposition 1 that are closely related to the credibility proportion of Hommes and Lustenhouwer (2019). First, the posterior collapses to a logistic function of the difference in performance between hypotheses, which is why credibility takes the discrete-choice form as in their case. Second, the sensitivity parameter, b, contrary to their free "intensity of choice", here is the precision of inflation innovations corrected for forgetting. Hence, when inflation is noisy (high σ_g^2), a given forecast error is weak evidence against either hypothesis, and credibility moves slowly, but when inflation is predictable (low σ_g^2), the same forecast error moves credibility fast.[11]

The next step is to calculate the predictive distribution of π_{t+1} weighted by credibility. I adopt the idea of averaging models that hold in different periods of time based on the "dynamic model averaging" of Koop and Korobilis (2012). Importantly, this approach incorporates credibility as part of the predictive distribution of each agent by itself and not as a proportion of the population as in Hommes and Lustenhouwer (2019). Therefore, the predictive distribution from averaging both hypotheses can be calculated, and directly its associated expectation, as Lemma 1 states.

[10] Following Koop and Korobilis (2012), only the errors observed up to time $t-1$ are considered when forecasting at time t.

[11] Since σ_g^2 is a perceived variance and not an estimated one, the calibration of Section 4 imports $b = 0.25$ from Hommes and Lustenhouwer (2019) rather than computing it from inflation data.

LEMMA 1 (Predictive distribution and the posterior-mean anchor). *Under Assumptions 1–4, cohort g's predictive distribution for π_{t+1} is the mixture*

$$\begin{aligned} f(\pi_{t+1} \mid I_t) &= \chi_{g,t}\,\varphi\big(\pi_{t+1};\,\mu_{T,t},\,\sigma_g^2\big) + (1-\chi_{g,t})\,\varphi\big(\pi_{t+1};\,\mu_{E,t},\,\sigma_g^2\big), \\ \mu_{h,t} &\equiv \rho_g \pi_t + (1-\rho_g) m_h, \end{aligned} \tag{10}$$

where $\varphi(x;\mu,\sigma^2)$ is the normal density, and its mean is

$$\begin{aligned} \mathbb{E}_{g,t}[\pi_{t+1}] &= \rho_g\,\pi_t + (1-\rho_g)\,m_{g,t}, \\ m_{g,t} &\equiv \chi_{g,t}\,\pi^* + (1-\chi_{g,t})\,m_{E,g} = \pi^* + (1-\chi_{g,t})(1-\nu)\,\Delta_g. \end{aligned} \tag{11}$$

where the distance between remembered inflation and the target,

$$\Delta_g \equiv \bar{\pi}_g - \pi^*,$$

is the cohort's experienced gap.

PROOF. See Appendix A, subsection A.2. □

As Equation (11) shows, the inflation expectation is the perceived law with the unknown anchor replaced by its posterior mean $m_{g,t}$, which is the target plus the part of the experienced gap that has not been removed by either credibility or communication, and which is ultimately where the central bank can exert its influence. Inflation expectations inherit the uncertainty about the anchor: when $\chi_{g,t} = 1$ (the *perfect-credibility scenario*), the anchor collapses to the target π^* regardless of the communication intensity selected by the central bank. At the opposite extreme, when $\chi_{g,t} = 0$ (the *null-credibility scenario*), the anchor collapses to $m_{E,g}$: expectations are anchored by the cohort's experiential inflation rate modulated by central-bank communication.

The ultimate step involves distorting the credibility-weighted predictive distribution by DE. Section 2 documents that household forecasts under-react to inflation news. Following Bordalo et al. (2018), I model this by assuming that the cohort does not report the mean of its predictive distribution but its *diagnostic* version, which reweights each state by how much more likely it has become since last period. Two restrictions determine where the distortion acts. First, in line with the DE literature, Assumption 5 says that it acts on the forecast and takes credibility as a state that is already known when the forecast is made. That is, inside the agent's mind, she first sets her subjective credibility and then reports her

diagnostic distribution. Otherwise the same $\theta < 0$ that produces under-reaction in the forecast would make credibility under-react too, and the credibility process would not be independently identified.

ASSUMPTION 5 (Diagnostic Distribution). *Cohort g does not report the mean of its predictive distribution. Instead, it reports its* diagnostic *version, the distribution*

$$f^{\theta}(\pi_{t+1} \mid I_t) \;\propto\; f(\pi_{t+1} \mid I_t)\left[\frac{f(\pi_{t+1} \mid I_t)}{f(\pi_{t+1} \mid I_{t-1})}\right]^{\theta}. \tag{12}$$

which reweights each value of π_{t+1} by how much more likely it has become since last period.

The parameter θ is common across cohorts. With $\theta > 0$ the cohort overweights states that became more likely (over-reaction), while $\theta < 0$ means it underweights them (under-reaction). Second, Assumption 6 says that DE does not act on the learning process, so hypotheses are evaluated by their own raw errors and not by the errors of the distorted forecast.

ASSUMPTION 6 (Learning Independent of DE). *The two hypotheses, $\{T, E\}$, are compared by the errors of their own forecasts under Equation* (3)*, not by the error of the diagnostic forecast, so the learning probabilities are not distorted.*

Following L'Huillier et al. (2024), direct application of the resulting diagnostic distribution of Equation (12) leads to:

$$\mathbb{E}_t^{\theta}[\pi_{t+1}] \;=\; \mathbb{E}_t[\pi_{t+1}] + \theta\big(\mathbb{E}_t[\pi_{t+1}] - \mathbb{E}_{t-1}[\pi_{t+1}]\big). \tag{13}$$

Proposition 2 linearises this equation around the steady state and delivers the inflation expectation, with endogenous credibility and under-reaction, that the next sections adopt throughout.

PROPOSITION 2 (Diagnostic Expectation with Endogenous Credibility). *For cohort g, the inflation expectation linearised around the steady state $\pi_t = \pi^*$ is*

$$\mathbb{E}_{g,t}^{\theta}[\hat{\pi}_{t+1}] \;=\; \rho_g\,\hat{\pi}_t + (1-\rho_g)\,\kappa_g\,\hat{\chi}_{g,t} + \theta\,\rho_g\big(\hat{\pi}_t - \rho_g\,\hat{\pi}_{t-1} - (1-\rho_g)\,\kappa_g\,\hat{\chi}_{g,t-1}\big). \tag{14}$$

where κ_g is the credibility sensitivity of the anchor,

$$\kappa_g \;=\; -(1-\nu)\,\Delta_g. \tag{15}$$

PROOF. See Appendix A, subsection A.3. □

The persistence ρ_g captures the idea that a cohort that lived through persistent inflation expects any inflationary shock to last longer, and revises its expectations more. On the other hand, the credibility sensitivity κ_g shows that cohorts whose experienced inflation is close to the target respond only slightly to credibility changes, whereas cohorts whose inflation experience is far from the target exhibit larger expectation adjustments when credibility changes. Since κ_g is negative for the cohorts that lived with inflation above the target, a fall in credibility raises their expected inflation. Finally, the last term is the diagnostic distortion. With $\theta < 0$ only a fraction $1+\theta$ of the revision of the forecast is incorporated, and the rest stays at last period's forecast. Finally, these three will be important when analyzing CB communication in Section 5, since the persistence and the diagnostic parameter will determine the diagnostic error of each cohort, and the credibility sensitivity will set its credibility error.

4. Credibility-Distorted Beliefs and the New-Keynesian Model

This section embeds the credibility-distorted beliefs of Section 3 into a Representative Agent New-Keynesian (RANK) model. The belief distortion essentially operates through the cohort's perceived real interest rate and the aggregate Euler equation, which ultimately causes cohorts to deviate from the FIRE intertemporal consumption path. Next, I extend the analysis to a richer model of the economy by considering a Heterogeneous-Agent New-Keynesian (HANK) model in which cohorts also differ in their marginal propensities to consume (MPC). This step makes clear how distorted beliefs now interact with the redistribution channel of monetary policy. While the *perceived real rate* affects consumption by distorting the intertemporal consumption path as in the RANK model, the *realised real rate*, the interest rate that households actually earn on their bonds, revalues each cohort's balance sheet through the URE channel. As Section 5 shows, the incidence in consumption shifts from the cohorts that forecast differently, as in the RANK model, to the cohorts that spend differently, as in the HANK model.

4.1. The Distorted RANK Model

This simple model consists of a discrete-time New Keynesian economy populated by a continuum of households divided into G cohorts. Every group is indexed by $g \in \{1, \ldots, G\}$ and has population weight $\omega_g \geq 0$, where $\sum_g \omega_g = 1$. Each cohort has the same preferences over consumption, $u(C) = C^{1-\sigma^{-1}}/(1-\sigma^{-1})$, and the same disutility from labor, $\zeta(N) = N^{1+\psi^{-1}}/(1+\psi^{-1})$, and receives the common income $y_{g,t} = y_t$, so there is no income heterogeneity. $N_{g,t}$ plays no role in the consumption-saving decision studied below. Households have access to a one-period nominal bond a_t paying the nominal interest rate i_t. Aggregate variables are expressed in uppercase letters while lowercase letters are reserved to deviations from the flexible-price steady state. Finally, the goods market clearing condition imposes $y_t = \int c_{i,t} \mathrm{d}i = \sum_g \omega_g c_{g,t}$, with $c_{g,t} \equiv \int_{i \in g} \frac{1}{\omega_g} c_{i,t} \mathrm{d}i$.

To cleanly isolate how the distorted-beliefs channel shapes consumption, I assume that the beliefs about future output of cohort g are FIRE, $\mathbb{E}_{g,t}[y_{t+h}] = y_{t+h}$, so distorted beliefs operate only through inflation expectations and hence through the perceived real rate. Distorted expectations may well affect how agents forecast output, but the BoE Inflation Attitudes Survey has no counterpart for output expectations, so imposing a parallel channel on y would mean accepting a mechanism that Section 2 cannot document, which is out of the scope of this paper.

Each cohort g solves the maximization problem:

$$\max_{\{C_{g,t}, a_{g,t+1}, N_{g,t}\}_{t \geq 0}} \mathbb{E}^{\theta}_{g,0} \sum_{t=0}^{\infty} \beta_g^t \left\{ u(C_{g,t}) - \zeta(N_{g,t}) \right\} \tag{16}$$

subject to

$$C_{g,t} + a_{g,t+1} = Y_t + \left(1 + r_t\right) a_{g,t}, \tag{17}$$

and the no-Ponzi condition $\lim_{T \to \infty} \beta_g^T a_{g,T} \geq 0$ holds, which binds with equality at the optimum (transversality).

Importantly, Proposition 2 implies that for every horizon $h \geq 0$, cohort g's perceived ex-ante real rate is

$$r^{\text{per}}_{g,t+h} = i_{t+h} - \rho_g \pi_{t+h} - (1-\rho_g) m_{g,t+h} - \theta \rho_g \left[u_{t+h} - (1-\rho_g) m_{g,t+h-1} \right]. \tag{18}$$

In this simple version of the economy, absent of wealth accumulation, the

realised real rate, r_t, only affects the return of current assets and not intertemporal consumption decisions, which depend on the perceived rate $r^{per}_{g,t}$ alone. Moreover, since cohorts hold no assets in steady state, $a^{ss}_g = 0$, there is no stock of wealth to be revalued when r_t moves, so the realised rate has no effect on consumption. This is the key difference with the HANK model, where there are heterogeneous asset positions, so the same movement in r_t revalues their cash-on-hand, the URE channel.

Intertemporal consumption decisions are pinned down by the perceived real rate of cohort g. This rate depicts the cohort-specific channels through which the distortion enters into the consumption-saving decision. The first mechanism states that cohorts that lived through persistent inflation (high ρ_g) perceive a lower real rate, whereas low-persistence cohorts perceive a real rate closer to the nominal rate i_{t+h}. The second channel propagates central-bank credibility $\chi_{g,t+h}$ through $\kappa_g = -(1-\nu)\Delta_g$ of Equation (15). It asserts that cohorts whose experienced inflation has persistently diverged from the target react more strongly to credibility changes. Finally, the diagnostic parameter θ measures how the cohort reacts to news. When $\theta < 0$, the cohort under-reacts to news, so the perceived real rate falls by less than the inflation news would demand, and households bring consumption forward by less than the FIRE case would justify.

The standard solution of the household maximization problem leads to the cohort-g distorted Euler equation:

$$C_{g,t}^{-1/\sigma} \;=\; \beta_g(1 + r^{\text{per}}_{g,t})\,\mathbb{E}^{\theta}_{g,t}[C_{g,t+1}^{-1/\sigma}] \tag{19}$$

It is worth noting that β_g introduces cohort heterogeneity, so the differences across cohorts in the perceived real interest rate prevail in steady state, $r^{per\,ss}_g = \beta_g^{-1} - 1$. The linearized version around the steady state for the perceived rate is

$$\hat{r}^{\text{per}}_{g,t+h} \;=\; \hat{\imath}_{t+h} - \rho_g\,\hat{\pi}_{t+h} - (1-\rho_g)\,\kappa_g\,\hat{\chi}_{g,t+h} - \theta\,\rho_g\big(\hat{\pi}_{t+h} - \rho_g\,\hat{\pi}_{t+h-1} - (1-\rho_g)\,\kappa_g\,\hat{\chi}_{g,t+h-1}\big). \tag{20}$$

while, Lemma 2 shows the equivalent expression for consumption:

LEMMA 2 (The cohort-g log-linearized consumption). *The first-order log-linearised solution of* (16)–(17) *around* $a^{ss}_g = 0$, $\pi^{ss} = \pi^*$, χ^{ss}_g, *and* $r^{per\,ss}_g = \beta_g^{-1} - 1$ *is*

$$c_{g,t} \;=\; (1-\beta_g)\sum_{h=0}^{\infty}\beta_g^h\,y_{t+h} \;-\; \sigma\sum_{h=0}^{\infty}\beta_g^{h+1}\,\hat{r}^{per}_{g,t+h}, \tag{21}$$

where $\hat{r}^{per}_{g,t+h}$ *is given by* (20).

PROOF. See Appendix A, subsection A.4. □

As a final step, assume that all cohorts smooth consumption at roughly the same speed, which implies $\beta_g^h \approx \bar{\beta}^h$ for all g, h, and define $\bar{\beta} \equiv \sum_g \omega_g \beta_g$, the population-average perceived rate $\bar{r}^{\text{per}}_t \equiv \sum_g \omega_g \hat{r}^{\text{per}}_{g,t}$, and recall the goods-market clearing condition $y_t = \sum_g \omega_g c_{g,t}$. Then, Equation 21 may be simplified to:

$$y_t \;=\; (1-\bar{\beta}) \sum_{h=0}^{\infty} \bar{\beta}^h \, y_{t+h} \;-\; \sigma \sum_{h=0}^{\infty} \bar{\beta}^{h+1} \, \bar{r}^{\text{per}}_{t+h}. \tag{22}$$

Cancelling common terms ultimately leads to:

$$y_t \;=\; y_{t+1} \;-\; \sigma \, \bar{r}^{\text{per}}_t.$$

Thus, the aggregation of the cohort-g Euler equations reduces to the representative-agent Euler equation.

To close the model, firms set prices à la Calvo and have perfect-foresight which yields the standard New Keynesian Phillips curve, augmented with an exogenous cost-push shock ε^{CP}_t that follows an AR(1) process with persistence ρ. The central bank conducts monetary policy through a Taylor rule that responds to inflation and is subject to an exogenous shock ε_t. Together with the aggregate Euler equation derived previously, the equilibrium of this simple economy is characterized by the following three-equation block:

$$\text{Policy Rule:} \quad i_t = \phi \pi_t + \varepsilon_t, \tag{23}$$

$$\text{NKPC:} \quad \pi_t = \varpi \, y_t + \beta \, \mathbb{E}_t[\pi_{t+1}] + \varepsilon^{CP}_t, \tag{24}$$

$$\text{Aggregate Euler:} \quad y_t = \mathbb{E}_t[y_{t+1}] - \sigma \, \bar{r}^{\text{per}}_t, \qquad \bar{r}^{\text{per}}_t \equiv \sum_g \omega_g \, \hat{r}^{\text{per}}_{g,t}. \tag{25}$$

The three-equation block (23)–(25), together with the distorted expectations of Proposition 2, is the system that the RANK simulations solve. Importantly for the welfare analysis developed later in the paper, cohorts react more or less strongly than the average, depending on their own distortion parameters ρ_g and κ_g, through their perceived real interest rate. To bring the model closer to the HANK literature, the next step is to introduce uninsurable income risk and heterogeneous asset positions, so that cohorts would differ in their marginal propensities to consume

(MPCs) and in the exposure (URE) of their balance sheets to r_t.

4.2. The Distorted HANK Model

This model builds on the HANK of Auclert et al. (2024), augmented with the distorted expectations mechanism introduced in Section 3 and the cohort belief heterogeneity structure of Guerreiro (2023). The economy consists of a continuum of infinitely-lived households $i \in [0,1]$, each belonging to one of G age cohorts that act as permanent types. Thus, cohorts do not age, and differ in their population weight ω_g, their discount factor $\beta_g \in (0,1)$ and their inflation history, $\bar{\pi_g}$.

As standard in the HANK literature, I assume that households face uninsurable idiosyncratic labour-productivity risk $e_{i,t}$, which follows a Markov chain with transition matrix $P(e, e')$. This matrix is common across cohorts and has stationary distribution $\pi^e(e)$ such that $\sum_e \pi^e(e)\, e = 1$.

Every household in group g receives a post-tax labour income from supplying $N_{g,t}$ hours of labor (determined by unions) given by:

$$y_{g,t}(e) \;=\; (1-\tau_t)\, e\, w_t\, N_{g,t}, \tag{26}$$

where w_t is the real wage and τ_t the labour-income tax. Finally, markets are incomplete and households trade a one-period real bond a such that $a \geq 0$.

To keep consistency with Proposition 2, I introduce the diagnostic-expectations mechanism *only* through distorted beliefs about the aggregate inflation path, which in turn shape the perceived real interest rate that drives consumption decisions through the *Fisher channel*. By contrast, I assume that households observe their idiosyncratic state $e_{i,t}$ and are fully rational with respect to it. Moreover, they know the transition matrix $P(e, e')$ and form correct expectations about their own income. Accordingly, the cohort-g perceived ex-ante real rate is,

$$r^{\text{per}}_{g,t} \;\equiv\; i_t - \mathbb{E}^{\theta}_{g,t}[\pi_{t+1}]. \tag{27}$$

Importantly, distorted beliefs operate only at the cohort level. The real return that wealth actually earns is the *realised* ex-post real rate, $r_{t+1} = i_t - \pi_{t+1}$. Thus, a bond's realised return does not depend on what its holder believes, so every household earns r_{t+1} on its bonds *regardless of its expectations*. The cohort-perceived rate $r^{\text{per}}_{g,t}$ in (27) is the belief about that same return held at the decision date, namely the Fisher equation evaluated under the diagnostic expectation $\mathbb{E}^{\theta}_{g,t}$, and it

is the rate that *shapes inter-temporal consumption decisions*. The two rates therefore coexist as a realised return and a forecast of it.

For every cohort-g household entering period t with bond holdings a and productivity e, the consumption and savings decisions are obtained from solving the Bellman equation

$$\begin{aligned} V_{g,t}(a,e) \;=\; \max_{c,\,a'\geq 0} \; u(c) \;+\; \beta_g\, \mathbb{E}^{\theta}_{g,t}\!\left[\, V_{g,t+1}(a',e') \,\middle|\, e \,\right] \\ \text{s.t.} \quad c + a' \;=\; y_{g,t}(e) \;+\; \left(1 + r_t\right) a, \end{aligned} \tag{28}$$

where $u(c) = c^{1-\sigma^{-1}}/(1-\sigma^{-1})$. The solution to Problem (28) defines the cohort consumption and asset policy functions $c_{g,t}(a,e)$ and $a_{g,t}(a,e)$.

Note that the budget constraint in (28) uses r_t on existing wealth, while the optimal saving decision discounts future marginal utility at the cohort-perceived rate. This implies the distorted Euler equation:

$$u'(c_{g,t}) \;=\; \beta_g \left(1 + r^{\text{per}}_{g,t}\right) \mathbb{E}^{\theta}_{g,t}\!\left[\, u'(c_{g,t+1}) \,\right], \tag{29}$$

so the discount rate between t and $t+1$ uses the nominal rate and the diagnostic expectation of the decision date t, exactly as in the model of Section 4.1.

As discussed in Guerreiro (2023), Problem (28) imposes the law of iterated expectations at the individual level, $\mathbb{E}^{\theta}_{g,t}[\mathbb{E}^{\theta}_{g,t+1}[\cdot]] = \mathbb{E}^{\theta}_{g,t}[\cdot]$, so households behave as if they do not expect to revise their inflation beliefs. Since diagnostic expectations do not distort the idiosyncratic state, the cross-sectional distribution of beliefs is orthogonal to (a,e). This condition allows the behavioural Jacobians of Appendix B.3 to be computed from their FIRE counterparts.

Finally, denote by $D_{g,t}(a,e)$ the joint distribution of assets and productivity. Given the policy $a_{g,t}(\cdot)$ and the exogenous chain P, it evolves according to[12]

$$D_{g,t+1}(a',e') \;=\; \sum_{e} \int \mathbf{1}\!\left\{ a_{g,t}(a,e) = a' \right\} P(e,e')\; D_{g,t}(\mathrm{d}a, e). \tag{30}$$

Hence, cohort aggregates may be calculated as

$$C_{g,t} \;=\; \sum_{e} \int c_{g,t}(a,e)\, D_{g,t}(\mathrm{d}a,e), \qquad A_{g,t} \;=\; \sum_{e} \int a_{g,t}(a,e)\, D_{g,t}(\mathrm{d}a,e), \tag{31}$$

[12] I implement the lottery method described in Young (2010) on a grid, so (30) becomes $D_{g,t+1} = \Lambda'_{g,t}\, D_{g,t}$ for a sparse transition matrix $\Lambda_{g,t}$.

while aggregate consumption and bond demand are given by $C_t = \sum_g \omega_g C_{g,t}$ and $A_t = \sum_g \omega_g A_{g,t}$.

The rest of the model is standard. The supply side of the economy, common across cohorts, follows Guerreiro (2023) and is the canonical sticky-wage New Keynesian block of Auclert et al. (2024). A competitive final-good firm operates the linear technology $Y_t = \Theta_t N_t$, so the real wage equals marginal productivity, $w_t = \Theta_t$. Price inflation is therefore wage inflation net of productivity growth. As in Guerreiro (2023), a continuum of unions aggregates cohort labour into N_t, sets wages subject to a quadratic adjustment cost, and delivers the linearised wage Phillips curve

$$\pi_t^w \;=\; \kappa_w\Big[\sigma^{-1}\hat{C}_t + \psi^{-1}\hat{N}_t - \big(\hat{w}_t - \hat{\tau}_t\big)\Big] \;+\; \bar{\beta}\,\mathbb{E}_t[\pi_{t+1}^w] \;+\; \varepsilon_t^{CP}, \tag{32}$$

with κ_w the slope, ψ the Frisch elasticity, ε_t^{CP} an exogenous cost-push shock, and hats denoting log-deviations. Cohort hours are allocated by the uniform union rule $N_{g,t} = \Gamma_g(N_t)$, with $\sum_g \omega_g \Gamma_g(N_t) = N_t$. Unions set wages and ration hours uniformly across households, so households do not choose labour directly and N does not enter their intertemporal problem. The government issues real debt $B_t = B$, spends $G_t = G$, and adjusts τ_t to satisfy $G + (1 + r_t)B = \tau_t Y_t + B$, and monetary policy follows a Taylor rule with an exogenous monetary shock ε_t^{MP},

$$i_t = r_{ss} + \phi_\pi \pi_t + \varepsilon_t^{MP}, \qquad r_{t+1} = i_t - \pi_{t+1}. \tag{33}$$

Finally, the equilibrium for this economy is defined in Appendix B.

4.2.1. Sequence-Space Representation

This paper uses the sequence space approach introduced by Auclert et al. (2021), in which the economy is written as a set of blocks and the first-order behaviour of a block with output O and input X is summarised by its Jacobian $\mathcal{J}^{O,X}$, whose (t, s) entry is the response of the date-t output to news about the date-s input. The only non-trivial blocks are the G cohort household blocks, whose Jacobians are computed with the fake-news algorithm. Appendix B.2 describes the remaining blocks and the solution method.

Through the functional $\mathcal{C}_g(\cdot)$, cohort g takes as inputs the realised real rate r_t, the perceived real rate $r_{g,t}^{\text{per}}$ of (27), and its after-tax labour income $(1 - \tau_t)\, w_t N_{g,t}$,

and returns as outputs cohort consumption $C_{g,t}$ and asset holdings $A_{g,t}$ with

$$\mathcal{C}_g\big(\{r_t, r^{\text{per}}_{g,t}, Y_{g,t}, \tau_t\}_{t\geq 0}\big) \;=\; \sum_e \int c_{g,t}(a,e)\, D_{g,t}(\mathrm{d}a, e), \tag{34}$$

and analogously for asset demand $\mathcal{A}_g$. The same block returns each cohort's consumption, with $\sum_g \omega_g \mathcal{C}_g = C$, used to calculate the cohort-level wedges for the HANK model in the next section.

The first-order response of cohort consumption $\mathbf{C}_g$ is a linear combination of its inputs' Jacobian contributions.[13]

$$\mathrm{d}\mathbf{C}_g \;=\; \mathcal{J}_g^{C,r} \cdot \mathrm{d}\mathbf{r} \;+\; \mathcal{J}_g^{C,r^{\text{per}}} \cdot \mathrm{d}\mathbf{r}^{\text{per}}_g \;+\; \mathcal{J}_g^{C,Y} \cdot \mathrm{d}\mathbf{Y}_g \;+\; \mathcal{J}_g^{C,\tau} \cdot \mathrm{d}\boldsymbol{\tau}. \tag{35}$$

The Jacobian $\mathcal{J}_g^{C,r}$ is the cash-on-hand response of cohort consumption to the realised real rate. Since bonds are in positive net supply and $a \geq 0$, every cohort is a net saver and a rise in r_t raises the cash-on-hand of all of them, by more for the wealthier old cohorts. Distorted beliefs introduce an additional mechanism through which the central bank's decisions may influence consumption. Expanding $\mathrm{d}\mathbf{r}^{\text{per}}_g$ in Equation 35 gives:

$$\begin{aligned}\mathcal{J}_g^{C,r^{\text{per}}}\mathrm{d}\mathbf{r}^{\text{per}}_g \;=\; &\underbrace{\mathcal{J}_g^{C,r^{\text{per}}}\mathrm{d}\mathbf{i}}_{\text{Fisher}} \;\underbrace{-\,\rho_g\,\mathcal{J}_g^{C,r^{\text{per}}}\big[(1+\theta)\mathbb{I} - \theta\rho_g L\big]\mathrm{d}\boldsymbol{\pi}}_{\text{Inflation}} \\ &\underbrace{-\,(1-\rho_g)\,\kappa_g\,\mathcal{J}_g^{C,r^{\text{per}}}\big[\mathbb{I} - \theta\rho_g L\big]\mathrm{d}\boldsymbol{\chi}_g}_{\text{Credibility}}.\end{aligned} \tag{36}$$

Equation 36 shows that, with L the lag operator, the perceived-rate Jacobian decomposes into a pure Fisher channel, an inflation component $\rho_g[(1+\theta)\mathbb{I} - \theta\rho_g L]\mathcal{J}_g^{C,r^{\text{per}}}$, and a credibility component $(1-\rho_g)\kappa_g[\mathbb{I} - \theta\rho_g L]\mathcal{J}_g^{C,r^{\text{per}}}$. The last one is the only channel that central-bank communication can influence, since its communication ν enters only through $m_{E,g}$.

To close the model, note that the cohort household blocks, the supply block (the Phillips curve), and the policy block jointly determine the aggregate paths $\mathbf{U} = (\mathbf{Y}, \boldsymbol{\pi}, \boldsymbol{\chi})'$ through the equilibrium conditions (32)–(A25). The model is solved in sequence space, taking output, inflation and credibility as unknowns, and imposing the asset-market clearing, the Phillips curve and the credibility law of

[13] In terms of notation, bold letters denote the stacked column vector of deviations at the different horizons, $\mathbf{X} = (X_0, X_1, \dots)'$.

motion conditions,

$$H(\mathbf{U}, \mathbf{Z}) = \mathbf{0}, \qquad \mathbf{Z} = (\varepsilon^{CP})', \tag{37}$$

where $\mathbf{Z}$ collects the cost-push shock, ε^{CP}. The welfare analysis of Section 5 considers only this shock, so the monetary shock ε_t^{MP} of (33) is shut at zero. Linearising around the steady state, the response of the aggregates to any shock is

$$\mathrm{d}\mathbf{U} = -H_{\mathbf{U}}^{-1} H_{\mathbf{Z}}\, \mathrm{d}\mathbf{Z}, \tag{38}$$

where $H_{\mathbf{U}}$ and $H_{\mathbf{Z}}$ are the corresponding Jacobians. The belief distortion enters this system only through the cohort household blocks, via their perceived real rate $r_{g,t}^{\text{per}}$ of (27). Therefore, a single Jacobian, $\mathcal{J}_g^{C,r^{\text{per}}}$, is enough to carry the inflation and the credibility components of (36) to consumption. Hence, no additional belief Jacobian or forecast-revision term is required, since all learning dynamics may be obtained from the credibility path χ_g. This is the representation that Proposition 4 uses in the next section to compute the cohort consumption wedges.

4.3. Calibration

Table 1 reports the calibration. Panel A collects the structural parameters of the three-equation block (23)–(25), set to the standard values of Auclert et al. (2024) and Galí (2015), together with the diagnostic parameter θ, imported from Carroll (2003). Since the DE operator imposes $\mathbb{E}^{\theta}_{g,t}[\pi_{t+1}] = \mathbb{E}_{g,t-1}[\pi_{t+1}] + (1+\theta)\big(\mathbb{E}_{g,t}[\pi_{t+1}] - \mathbb{E}_{g,t-1}[\pi_{t+1}]\big)$, the diagnostic expectation incorporates a fraction $1+\theta$ of the forecast revision on impact, which matches the share $\lambda = 0.27$ of households that update each quarter in Carroll (2003), so that $\theta = \lambda - 1 = -0.73$. The learning parameters ρ_U and b complete the panel. Both are imported from Hommes and Lustenhouwer (2019): the memory of past forecast errors is $\rho_U = 0.5$ and the learning intensity is $b = 0.25$. Panel B builds the cohort moments from UK data. The demographic weights ω_g are UK ONS mid-2019 shares, while the lifetime inflation persistence ρ_g and experienced mean $\bar{\pi}_g$ are constructed from the ONS quarterly-annualised CPI series (Appendix C.1), measured at 2019Q4 over the window ending in 2019Q3. Two features of Panel B drive the results. The cohorts under 55 lived through inflation that was close to the target and barely persistent, while the cohorts over 55 lived through inflation that was two to three points above the target and much more persistent. The first difference determines the credibility error of each cohort and

the second its diagnostic error.

TABLE 1. **Calibration of the simple model.**

Panel A. Structural and belief parameters

Symbol	Value	Description	Source
β	0.99	Discount factor, quarterly	Auclert et al. (2024)
σ	1.00	Intertemporal elasticity of substitution	Auclert et al. (2024)
ϖ	0.03	New Keynesian Phillips curve slope	Auclert et al. (2024)
ϕ	1.50	Taylor-rule coefficient on inflation	Auclert et al. (2024)
ρ	0.50	AR(1) persistence of the shocks ε_t, ε_t^{CP}	Galí (2015)
π^*	2.00%	Inflation target	
θ	–0.73	Diagnostic parameter	Carroll (2003)
ρ_U	0.50	Memory of the discounted likelihood	Hommes and Lustenhouwer (2019)
b	0.25	Learning intensity	Hommes and Lustenhouwer (2019)
ν	0.00	Intensity of target-communication	
T	40	Horizon for the IRFs (quarters)	

Panel B. Cohort-level calibration

	15–24	25–34	35–44	45–54	55–64	65+	Pop. mean
Population weight, ω_g	0.148	0.163	0.153	0.168	0.146	0.222	1.000
Inflation persistence, ρ_g	0.113	0.106	0.035	0.179	0.559	0.554	0.274
Experienced inflation, $\bar{\pi}_g$ (%)	2.02	2.35	2.00	2.66	4.21	5.04	3.17

Notes: ρ_g and $\bar{\pi}_g$ are the AR(1) slope and the mean of the cohort's lifetime quarterly-annualised UK CPI inflation up to 2019Q4. The last column of Panel B reports population-weighted means.

The HANK model shares Panel A and Panel B and shows the cohort discount factors. Each cohort's β_g is set so that the model reproduces the cohort wealth-to-income ratios built from the ONS Wealth and Assets Survey (mean total wealth per adult, 2018–20) and the BoE IAS income bands (waves up to 2021Q4), with the average discount factor pinned down by the asset-market clearing condition. Within each cohort, a fraction of households carries a lower discount factor, which generates the hand-to-mouth households of the model, as in Auclert et al. (2025). The marginal propensities to consume are an outcome of this calibration and fall with age and wealth. Appendix D.3 reports the data, the procedure and the resulting β_g.

5. Welfare and Central Bank Communication

This section makes the case that central bank communication has welfare and distributional effects when beliefs are distorted. In the RANK model with FIRE, every cohort forecasts future inflation correctly and picks the consumption path it would choose with perfect foresight, so there is no welfare loss from beliefs. However, under distorted beliefs this is no longer true. Each cohort's experienced inflation and credibility shape its forecast of inflation and its perceived real interest rate. This heterogeneity in Euler equations is what ultimately changes each cohort's consumption path, since each cohort departs from the path it would choose if it forecast inflation correctly. According to Gabaix (2020) this deviation is a valid welfare measure since it is evaluated with objective, rather than subjective expectations. In this sense, the cost associated with a belief distortion is the utility lost along the consumption path it induces, relative to the path the household would choose with correct beliefs, which is why this paper adopts it as its measure of welfare throughout.

5.1. Consumption Wedges

Define the *rational path*, $dc^{rat}_{g,t}$, as the consumption path cohort g would choose for the same realised paths of prices and income, under *perfect foresight of inflation*. Therefore, the rational path measures the incidence of the shock with correct beliefs, and the gap between the actual path and the rational path measures the cost of distorted beliefs. However, the reference path behaves differently across models. In the RANK model the only thing that changes across cohorts is inflation expectations, so the rational path is common to all cohorts. On the other hand, in the HANK model the rational path is cohort specific, because even if a cohort has correct beliefs its consumption path responds to its idiosyncratic factors, such as the effect of the realised real rate through its asset position, and of its labour income and taxes through its MPC.

Belief heterogeneity makes each cohort depart from the consumption path it would choose with correct beliefs, and that deviation is costly. Define the cohort-g *consumption wedge*, in the RANK and in the HANK model respectively, as the difference between the consumption path cohort g chooses under its *own* distorted belief and its *rational path* dc^{rat}, common to all cohorts in the RANK and cohort specific in the HANK:

$$\mathcal{W}^{R}_{g,t} \;\equiv\; dc_{g,t} - dc^{rat}_{t}, \qquad \mathcal{W}^{H}_{g,t} \;\equiv\; dc_{g,t} - dc^{rat}_{g,t}. \tag{39}$$

Importantly, consumption wedges isolate the effect of distorted beliefs on consumption and they are zero when the cohort forecasts correctly. Once the equilibrium paths of inflation and credibility are known, the wedges follow immediately. Next, define the *cohort-g* and *aggregate* welfare losses of each model ($M \in \{R, H\}$), respectively, as:

$$L^{M}_{g,t} \;\equiv\; \tfrac{1}{2}\big(\mathcal{W}^{M}_{g,t}\big)^{2}, \qquad \mathcal{L}^{M}_{t} \;\equiv\; \sum_{g} \omega_{g}\, L^{M}_{g,t}, \qquad M \in \{R, H\}. \tag{40}$$

Proposition 3 characterises the consumption wedges and welfare losses in the RANK model, where only beliefs differ across cohorts, and Proposition 4 does so in the HANK model, where balance sheets and income differ as well.

PROPOSITION 3 (Consumption Wedges and Welfare in RANK). *Consider any sequence of structural shocks that produces equilibrium paths $\{\hat{\imath}_t, \hat{\pi}_t, \hat{\chi}_{g,t}\}_{t\geq 0}$ of the RANK model. Recall, the linearized inflation expectations follows $\mathrm{d}\mathbb{E}^{\theta}_{g,t}[\hat{\pi}_{t+1}] = \Lambda^{\pi}_{g}\hat{\pi}_t + \Lambda^{\chi}_{g}\hat{\chi}_{g,t}$ with*

$$\Lambda^{\pi}_{g}\hat{\pi}_t \;\equiv\; \rho_g\big[(1+\theta)\hat{\pi}_t - \theta\rho_g\hat{\pi}_{t-1}\big], \qquad \Lambda^{\chi}_{g}\hat{\chi}_{g,t} \;\equiv\; (1-\rho_g)\,\kappa_g\big[\hat{\chi}_{g,t} - \theta\rho_g\hat{\chi}_{g,t-1}\big]. \tag{41}$$

*Let $A_g \equiv -U''_g(c^{rat}_{g,t}) > 0$ denote cohort g's coefficient of absolute risk aversion at the rational path. Then, the **cohort-g consumption wedge**, $\mathcal{W}^{R}_{g,t} \equiv dc_{g,t} - dc^{rat}_{t}$, could be expressed as:*

$$\mathcal{W}^{R}_{g,t} \;=\; \sigma\sum_{s\geq 0}\Big(\mathrm{d}\mathbb{E}^{\theta}_{g,t+s}[\hat{\pi}_{t+s+1}] - \hat{\pi}_{t+s+1}\Big) \;=\; \sigma\sum_{s\geq 0}\Big[\big(\Lambda^{\pi}_{g}\hat{\pi}_{t+s} - \hat{\pi}_{t+s+1}\big) \;+\; \Lambda^{\chi}_{g}\hat{\chi}_{g,t+s}\Big]. \tag{42}$$

*Moreover, define the **diagnostic error** as $\mathcal{A}_{g,t}$ and the **credibility error** by $(1-\nu)\mathcal{B}_{g,t}$, where*

$$\begin{aligned} \mathcal{A}_{g,t} \;&\equiv\; \sigma\sum_{s\geq 0}\big(\Lambda^{\pi}_{g}\hat{\pi}_{t+s} - \hat{\pi}_{t+s+1}\big), \\ \mathcal{B}_{g,t} \;&\equiv\; -\sigma\sum_{s\geq 0}(1-\rho_g)\,\Delta_g\, X_{g,t+s}, \qquad X_{g,t} \;\equiv\; \hat{\chi}_{g,t} - \theta\rho_g\hat{\chi}_{g,t-1}, \end{aligned} \tag{43}$$

so that $\mathcal{W}^R_{g,t} = \mathcal{A}_{g,t} + (1-\nu)\mathcal{B}_{g,t}$. *Then, the* ***cohort-g welfare loss*** *is the utility the cohort loses along its chosen path relative to its rational path,*

$$L^R_{g,t} \;=\; \frac{U_g\big(c^{rat}_{g,t}\big) - U_g\big(c_{g,t}\big)}{A_g} \;=\; \tfrac{1}{2}\mathcal{A}^2_{g,t} \;+\; (1-\nu)\,\mathcal{A}_{g,t}\mathcal{B}_{g,t} \;+\; \tfrac{1}{2}(1-\nu)^2\mathcal{B}^2_{g,t}, \qquad (44)$$

while the ***aggregate welfare loss*** *of* (40) *is,*

$$\mathcal{L}^R_t \;=\; \sum_g \omega_g\Big[\tfrac{1}{2}\mathcal{A}^2_{g,t} \;+\; (1-\nu)\,\mathcal{A}_{g,t}\mathcal{B}_{g,t} \;+\; \tfrac{1}{2}(1-\nu)^2\mathcal{B}^2_{g,t}\Big]. \qquad (45)$$

PROOF. See Appendix A, subsection A.5. □

Proposition 3 states that both the cohort and the aggregate welfare losses depend only on the equilibrium paths $(\hat{\pi}_t, \hat{\chi}_{g,t})$, and that the wedge is the sum of two forecast errors. First, the *diagnostic error* is the gap between the cohort's inflation forecast with the credibility channel turned off and realised inflation. It is pinned down by Λ^{π}_g, which ultimately depends on the cohort's lifetime inflation persistence ρ_g and on θ. With under-reaction ($\theta < 0$) this error is negative on impact for all cohorts. They expect too little inflation, perceive a real rate that is too high, and postpone consumption relative to their rational path. For the cohorts under 55 it stays negative while the shock lasts. For the cohorts over 55 it turns positive after impact, because the persistence they lived through is higher than the persistence of inflation after the shock, and under-reaction keeps part of last period's forecast alive. Importantly, this channel is not affected by how the central bank decides to communicate with the public.

Second, the *credibility error* is the part of the forecast that modulates toward which anchor the cohort inclines, $\Lambda^{\chi}_g\hat{\chi}_{g,t}$. This channel is positive for the cohorts that have lived with inflation above the target, because they expect more inflation, and it is the only part of the consumption wedge that the central bank can move (even eliminate) with its communication via ν. Proposition 3 also shows the channels through which the communication of the central bank may shape the cohort loss. First, the contribution of the diagnostic error, $\mathcal{A}^2_{g,t}$. Second, the two terms that the central bank can eliminate with full-target communication, namely the contribution of the credibility error, $\mathcal{B}^2_{g,t}$, and the cross product of the two errors, $\mathcal{A}_{g,t}\mathcal{B}_{g,t}$.

Proposition 4 shows the consumption wedges of the HANK model. An important difference in this model is that the consumption path of cohort g additionally responds to the realised real rate through its asset position (the URE channel) and to income and taxes. However, these channels are the same under distorted beliefs and under the rational path, because both sequences face the same realised prices and income. This implies that they cancel when calculating the consumption wedge, leading the latter to depend only on the forecast error, discounted by the cohort's own consumption response.

PROPOSITION 4 (Consumption Wedges and Welfare in HANK). *Let* $\mathrm{d}\mathbf{C}_g$ *be the cohort-g consumption path generated by* (35)*, with the perceived-rate path in* (36)*. Define the cohort-g per-capita consumption path as* $\mathrm{d}c_{g,t} \equiv \mathrm{d}C_{g,t}/C_g$*, the Jacobians* $\mathcal{J}_g^{c,X} \equiv \mathcal{J}_g^{C,X}/C_g$*, and* $\mathcal{J}_g^{c,\pi} \equiv -\mathcal{J}_g^{c,r^{\text{per}}}$*, the consumption response to the cohort's subjective inflation forecast* $\mathrm{d}\mathcal{E}_g \equiv \Lambda_g^{\pi}\mathrm{d}\boldsymbol{\pi} + \Lambda_g^{\chi}\mathrm{d}\boldsymbol{\chi}_g$*, with* Λ_g^{π} *and* Λ_g^{χ} *the stacked operators of* (41)*. Let* $\mathrm{d}\mathbf{c}_g^{rat}$ *be the* ***rational path*** *of cohort g under perfect foresight inflation expectations. Then, the cohort consumption wedge* $\mathcal{W}_{g,t}^H \equiv \mathrm{d}c_{g,t} - \mathrm{d}c_{g,t}^{rat}$ *is the sum of the* ***diagnostic error*** *and a* ***credibility error*** *modulated by* $\mathcal{J}_g^{c,\pi}$*,*

$$\begin{aligned}\mathcal{W}_{g,t}^H &= \left[\mathcal{J}_g^{c,\pi}\left(\mathrm{d}\mathcal{E}_g - \mathbb{F}\,\mathrm{d}\boldsymbol{\pi}\right)\right]_t \\ &= \left[\mathcal{J}_g^{c,\pi}\left(\Lambda_g^{\pi} - \mathbb{F}\right)\mathrm{d}\boldsymbol{\pi}\right]_t + \left[\mathcal{J}_g^{c,\pi}\Lambda_g^{\chi}\,\mathrm{d}\boldsymbol{\chi}_g\right]_t.\end{aligned} \tag{46}$$

Denote the diagnostic error by $\mathcal{A}_{g,t}^H$ *and the credibility error by* $(1-\nu)\mathcal{B}_{g,t}^H$*,*

$$\begin{aligned}\mathcal{A}_{g,t}^H &\equiv \left[\mathcal{J}_g^{c,\pi}\left(\Lambda_g^{\pi} - \mathbb{F}\right)\mathrm{d}\boldsymbol{\pi}\right]_t, \\ \mathcal{B}_{g,t}^H &\equiv -(1-\rho_g)\,\Delta_g\left[\mathcal{J}_g^{c,\pi}\mathbf{X}_g\right]_t, \qquad \mathbf{X}_g \equiv (\mathbb{I} - \theta\rho_g L)\,\mathrm{d}\boldsymbol{\chi}_g,\end{aligned} \tag{47}$$

so that $\mathcal{W}_{g,t}^H = \mathcal{A}_{g,t}^H + (1-\nu)\mathcal{B}_{g,t}^H$*. As in the RANK model, the* ***cohort-g welfare loss*** *of* (40) *is*

$$L_{g,t}^H = \tfrac{1}{2}\left(\mathcal{A}_{g,t}^H\right)^2 + (1-\nu)\,\mathcal{A}_{g,t}^H\mathcal{B}_{g,t}^H + \tfrac{1}{2}(1-\nu)^2\left(\mathcal{B}_{g,t}^H\right)^2, \tag{48}$$

while the ***aggregate welfare loss*** *is:*

$$\mathcal{L}_t^H = \sum_g \omega_g\left[\tfrac{1}{2}\left(\mathcal{A}_{g,t}^H\right)^2 + (1-\nu)\,\mathcal{A}_{g,t}^H\mathcal{B}_{g,t}^H + \tfrac{1}{2}(1-\nu)^2\left(\mathcal{B}_{g,t}^H\right)^2\right]. \tag{49}$$

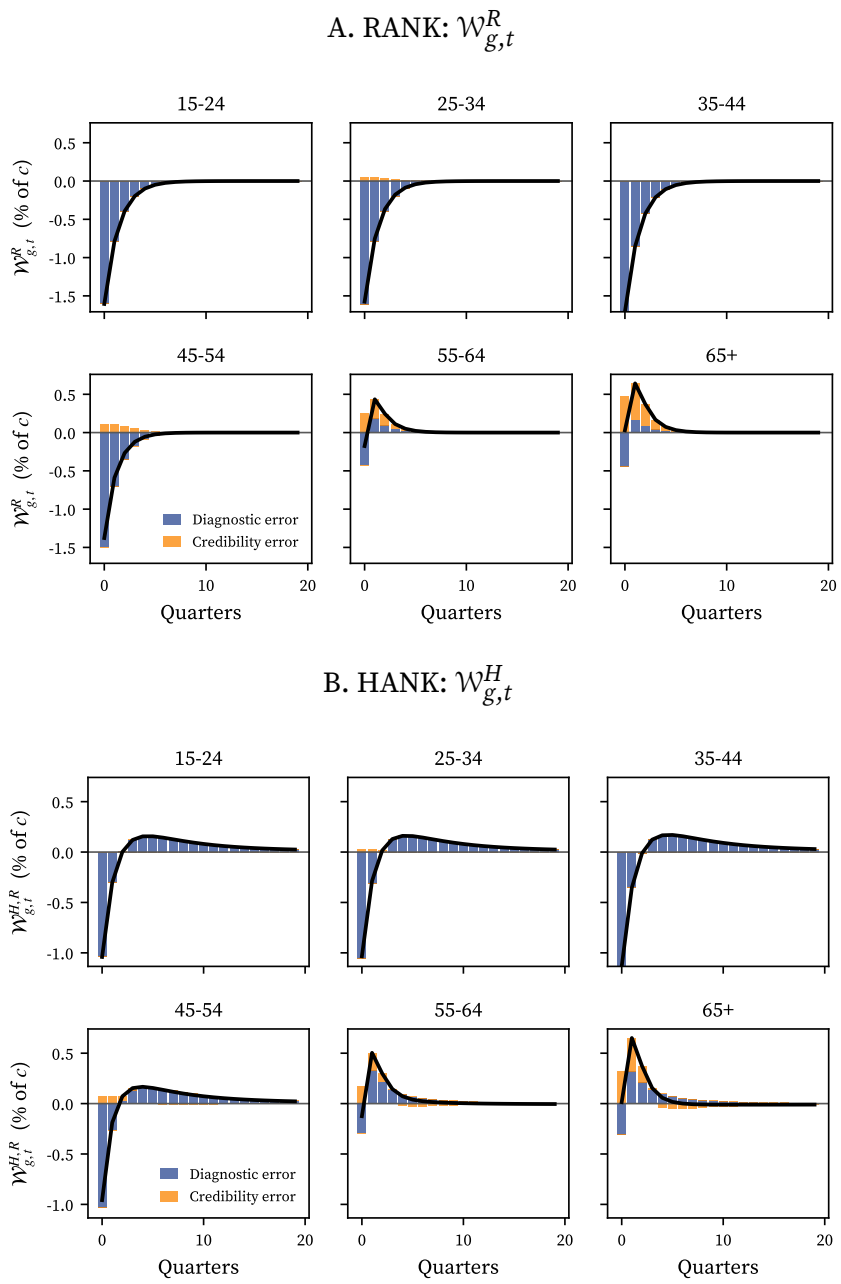
PROOF. See Appendix A, subsection A.6. □

Proposition 4 shows that the HANK wedge has the same two parts as the RANK wedge, a diagnostic error and a credibility error. However, the HANK model changes the wedge in one way. The diagnostic and credibility errors are scaled by the cohort's consumption Jacobian $\mathcal{J}_g^{c,\pi}$, the intertemporal substitution response to the perceived rate, which is smaller than the permanent-income response of the RANK model because the households close to the borrowing constraint do not substitute intertemporally. Hence the same forecast error moves consumption by less, and the response is spread over time rather than taken all at once. Balance sheets and income still matter through the rational path $dc_{g,t}^{rat}$. Since younger cohorts, which hold the smallest asset positions, depend most on labour income and have the highest MPC, they cut consumption the most when income falls after the shock.

Figure 5 shows the cohort wedges after the cost-push shock, and both models confirm the same story. For the younger cohorts (under 55) the diagnostic error of the wedge is negative, implying that they under-react to the shock, expect too little inflation and under-consume on impact relative to their rational path. Their credibility error is negligible, because their experienced inflation is close to the target. For the older cohorts (55 or above) the diagnostic error is negative only on impact because the persistent inflation they lived through makes them expect higher inflation. Additionally, their credibility error is consistently positive. On impact it offsets more than half of their under-reaction, and afterwards it adds to a diagnostic error that has turned positive. Since they have lived with inflation consistently above the target, they expect more inflation than realised, and over-consume. The HANK model has smaller wedges, because only unconstrained households substitute intertemporally, while the hand-to-mouth cannot.

The cumulated losses $L_g^M(0)$ rank the cohorts accordingly. The four cohorts under 55 account for 92% of the aggregate loss in the RANK model and for 82% in the HANK model while and the two cohorts over 55 for the remaining. The older cohorts lose less because their diagnostic error is smaller and somewhat their credibility error offsets the under-reaction. This is the part of the loss that communication can remove, and the next section examines it.

FIGURE 5. Cohort Consumption Wedges After a +1pp Cost-Push



5.2. Communication Design

This section asks which intensity of target-communication the central bank should choose and shows that the answer is not always full intensity. It then shows how the gain from communication is distributed across cohorts, how communication interacts with traditional monetary policy, and the cost of delaying it.

5.2.1. Optimal Target Communication

Propositions 3 and 4 show that only the credibility error reacts to ν, and that it is larger for the cohorts with a high inflationary experience, whose forecasts lean toward their own history when a cost-push shock hits. For these cohorts the credibility error is positive, while the diagnostic error is negative on impact due to under-reaction, so the two errors partly cancel each other out there. Their inflation history pushes their forecast up exactly when under-reaction holds it down. Hence, the credibility error of these cohorts works partly as a correction of their under-reaction and partly as an over-prediction. Communication removes the over-prediction, but it also removes the correction, and losing the correction is a communication loss.

To elaborate on this idea, first define $L_g^R(\nu) = \frac{1}{2}\sum_{t\geq 0}\big(\mathcal{W}_{g,t}^R(\nu)\big)^2$ as the cumulated welfare loss of cohort g at the intensity ν in the RANK model, with its identical counterpart $L_g^H(\nu) = \frac{1}{2}\sum_{t\geq 0}\big(\mathcal{W}_{g,t}^H(\nu)\big)^2$ for the HANK model. Next, define the *offsetting share*, $\mathcal{O}_g$, as

$$\mathcal{O}_g \equiv -\frac{\sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t}}{\sum_{t\geq 0}\mathcal{B}_{g,t}^2}, \tag{50}$$

The offsetting share is the fraction of the credibility error that works against the diagnostic error. When $\mathcal{O}_g$ is positive the credibility error offsets part of the under-reaction while when it is negative the two errors push consumption in the same direction. A similar expression for the offsetting share may be obtained for the HANK model, $\mathcal{O}_g^H$:

$$\mathcal{O}_g^H \equiv -\frac{\sum_{t\geq 0}\mathcal{A}_{g,t}^H\mathcal{B}_{g,t}^H}{\sum_{t\geq 0}\big(\mathcal{B}_{g,t}^H\big)^2}. \tag{51}$$

Using these definitions, in the RANK model the cohort loss may be expressed as:

$$L_g^R(\nu) = \frac{1}{2}\sum_{t\geq 0}\mathcal{A}_{g,t}^2 + \frac{1}{2}\Big[(1-\nu)^2 - 2(1-\nu)\,\mathcal{O}_g\Big]\sum_{t\geq 0}\mathcal{B}_{g,t}^2. \tag{52}$$

while in the HANK model as:

$$L_g^H(\nu) = \frac{1}{2}\sum_{t\geq 0}\big(\mathcal{A}_{g,t}^H\big)^2 + \frac{1}{2}\Big[(1-\nu)^2 - 2(1-\nu)\,\mathcal{O}_g^H\Big]\sum_{t\geq 0}\big(\mathcal{B}_{g,t}^H\big)^2. \tag{53}$$

While the central bank cannot reach the first term of Equation (52), since it

captures the effects of the persistence of lived inflation (ρ_g) or the diagnosticity parameter (θ), it may eliminate its remaining components, which operate through the credibility error. Proposition 5 then characterises the optimal choice of ν that minimises these components.

PROPOSITION 5 (Optimal Target Communication). *Consider a cost-push shock with realised, already observed, equilibrium paths $\{\hat{\pi}_t, \hat{\chi}_{g,t}\}_{t\geq 0}$ of inflation and cohort credibility. Let the cumulated losses in the RANK model be $L_g^R(\nu) = \frac{1}{2}\sum_{t\geq 0}\big(\mathcal{W}_{g,t}^R(\nu)\big)^2$ and $\mathcal{L}^R(\nu) = \sum_g \omega_g L_g^R(\nu)$. Then the intensity cohort g would choose and the aggregate optimum are*

$$\nu_g^{\star,R} \equiv \arg\min_{\nu\in\mathbb{R}} L_g^R(\nu) = 1 - \mathcal{O}_g,$$

$$\nu_{\text{agg}}^{\star,R} \equiv \arg\min_{\nu\in\mathbb{R}} \mathcal{L}^R(\nu) = 1 - \sum_g \mathcal{S}_g^R\, \mathcal{O}_g, \qquad \mathcal{S}_g^R \equiv \frac{\omega_g \sum_{t\geq 0} \mathcal{B}_{g,t}^2}{\sum_h \omega_h \sum_{t\geq 0} \mathcal{B}_{h,t}^2},$$

whenever the denominators are non-zero, where $\mathcal{S}_g^R$ is the share of cohort g in the population-weighted credibility error. A similar expression holds for the HANK model:

$$\nu_g^{\star,H} \equiv \arg\min_{\nu\in\mathbb{R}} L_g^H(\nu) = 1 - \mathcal{O}_g^H,$$

$$\nu_{\text{agg}}^{\star,H} \equiv \arg\min_{\nu\in\mathbb{R}} \sum_g \omega_g\, L_g^H(\nu) = 1 - \sum_g \mathcal{S}_g^H\, \mathcal{O}_g^H, \qquad \mathcal{S}_g^H \equiv \frac{\omega_g \sum_{t\geq 0}\big(\mathcal{B}_{g,t}^H\big)^2}{\sum_h \omega_h \sum_{t\geq 0}\big(\mathcal{B}_{h,t}^H\big)^2}.$$

whenever the denominators are non-zero.

PROOF. See Appendix A, subsection A.7. □

Proposition 5 shows that the intensity of communication each cohort chooses depends on its offsetting share, and Table 2 reports its values for the calibration. The cohorts whose experienced inflation is at the target (the 15–24 and the 35–44) have no credibility error, so their offsetting share is not identified and any intensity is optimal for them. The 25–34 and the 45–54, whose experienced inflation is only slightly above the target, have a small credibility error that offsets their diagnostic error. Hence, removing the over-prediction of inflation with intense communication only exposes their under-reaction, so they lose with this policy.

The two cohorts over 55, with the highest experienced inflation, are where the models differ. In the RANK model, their experienced inflation and the loss of

credibility it brings make them excessively over-predict inflation, and the under-reaction of diagnostic beliefs corrects only part of it. On the other hand, in the HANK model the offsetting share is negative, showing that after impact these cohorts in effect switch from under-predicting to over-predicting inflation, so that their diagnostic and credibility errors complement to a higher loss unless communication removes the credibility error. The reason is that high-MPC consumers close to the borrowing constraint cut consumption aggressively, so the contraction in demand is more persistent and inflation returns to the target faster, while the older cohorts, who forecast from their own inflation experience, keep over-predicting inflation for several periods. Therefore, anchoring them immediately with an intensive communication is optimal.

TABLE 2. **Optimal Target Communication Primitives**

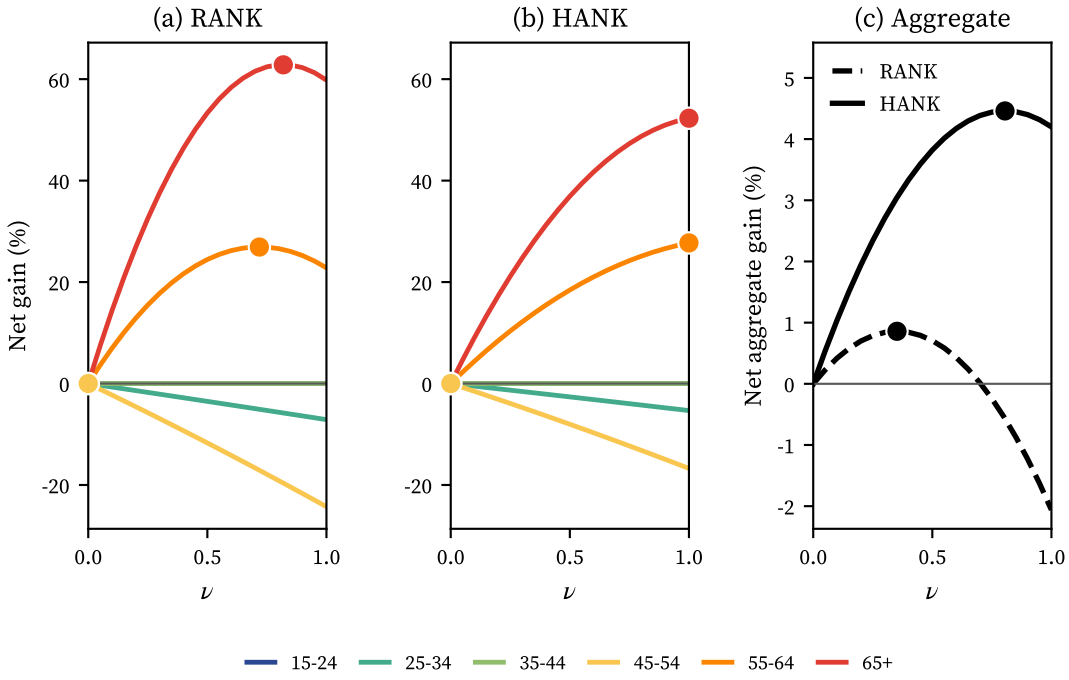
	RANK			HANK		
Cohort	$\mathcal{O}_g^R$	$\nu_g^{\star,R}$	$\mathcal{S}_g^R$	$\mathcal{O}_g^H$	$\nu_g^{\star,H}$	$\mathcal{S}_g^H$
15–24	–	–	–	–	–	–
25–34	23.55	−22.55	0.01	21.04	−20.04	0.00
35–44	–	–	–	–	–	–
45–54	7.59	−6.59	0.04	6.64	−5.64	0.04
55–64	0.28	0.72	0.15	−0.26	1.26	0.15
65+	0.18	0.82	0.80	−0.11	1.11	0.81
Aggregate	0.65	0.35	1	0.19	0.81	1

The optimum of Proposition 5 shows that each cohort demands a different intensity of communication, depending on its offsetting share, while the central bank can choose only one. The aggregate optimum weights the cohort optima, so it benefits the cohorts whose optimum lies close to or above it and is costly for the cohorts whose optimum lies far below it. To show this cost-benefit scenario, consider the case of no communication as the reference ($\nu = 0$). Define the *net communication gain* of cohort g from the intensity ν relative to $\nu = 0$,

$$\begin{aligned}\text{Net Gain}_g^M(\nu) &\equiv 100\left(\frac{L_g^M(0)-L_g^M(\nu)}{L_g^M(0)}\right)\\ &= 100\Bigg[\underbrace{\tfrac{1}{2}\nu(2-\nu)}_{\substack{\text{Benefit: credibility}\\ \text{error removed}}} - \underbrace{\nu\,\mathcal{O}_g^M}_{\substack{\text{Cost: offsetting}\\ \text{part lost}}}\Bigg]\frac{\sum_{t\geq 0}\left(\mathcal{B}_{g,t}^M\right)^2}{L_g^M(0)}, \qquad M\in\{R,H\}.\end{aligned} \tag{54}$$

Equation (54) shows the decreasing marginal benefit of target communication in eliminating part of the credibility error.[14] Hence, full intensity always provides a positive benefit, but communication may also have a cost, depending on the sign of $\mathcal{O}_g$.

FIGURE 6. **Net Gain from Target Communication**



Note: Dots show the cohort optima $\nu_g^{\star,R}$ and $\nu_g^{\star,H}$ of Proposition 5, at the corner when they lie outside $[0,1]$; in panel (c), the aggregate optima $\nu_{\text{agg}}^{\star,R}$ and $\nu_{\text{agg}}^{\star,H}$.

[14] The cohort optima $\nu_g^{\star,M}$ previously obtained are also the optima of this cost-benefit problem, since the net gain is a linear transformation of $L_g^M(\nu)$.

Figure 6 reports the net gain of each cohort. The two cohorts over 55 gain the most, since they hold the largest credibility error and the benefit of removing it exceeds the offsetting part they lose. Consistent with the previous discussion, in the RANK those cohorts optimally keep part of the credibility channel, while in the HANK both errors push consumption in the same direction and the central bank should remove the credibility error entirely. The 15–24 and the 35–44 have no credibility error, so their gain is zero at every intensity. Finally, the 25–34 and especially the 45–54 lose with communication at every intensity, for the reason already discussed.

Panel (c) aggregates the gains and shows the importance of the $\mathcal{S}_g^M$ of Proposition 5 in determining which cohorts pull the aggregate optimum. As Table 2 shows, the cohorts over 55 receive around 95% of this weight because they hold the largest credibility error, and on their own they would choose an intensity of about 0.8 in the RANK model and full communication in the HANK model. The remaining weight belongs to the 25–34 and the 45–54, whose offsetting share is so large that in the RANK model they bring the aggregate optimum down to 0.35. In the HANK model the switch of the older cohorts from under-predicting to over-predicting inflation dominates, and the aggregate optimum more than doubles, to 0.81. The aggregate gain is modest, about 1% of the loss in the RANK model and 4.5% in the HANK model, because communication only reaches the credibility error, but it is large for the cohorts over 55. This leaves the lesson that the effects of target communication are mainly distributional, and that the persistent over-prediction of the HANK model, inherent to its income heterogeneity, offers the central bank a normative reason to implement it intensively.

5.2.2. Communication, Monetary Policy and Timing

Once the optimal target communication is characterised, this paper examines two additional topics. First, how target communication interacts with the traditional interest-rate adjustments, analysed by Goy et al. (2022) for forward guidance under imperfect credibility and empirically by Coibion et al. (2022). Second, once the economy is hit by an inflationary shock, what the associated costs of delayed communication are, in the timing sense of Bomfim and Rudebusch (2000) for disinflation under imperfect credibility.

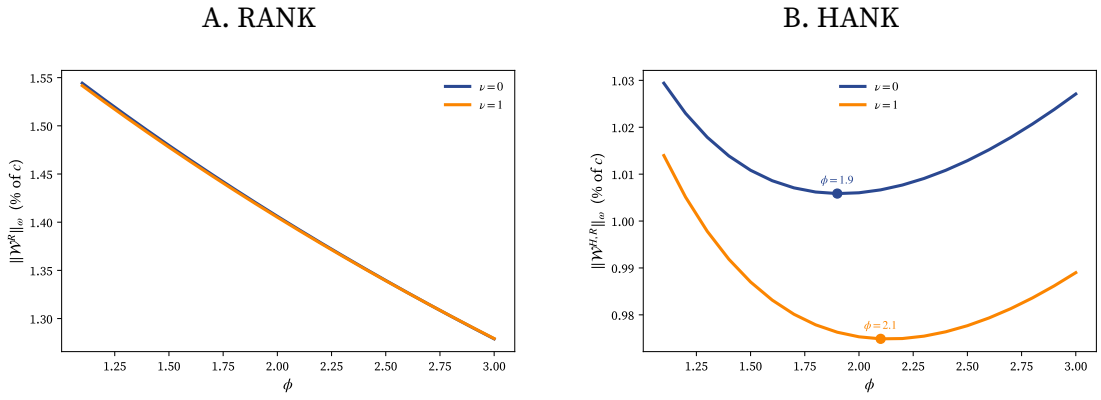
Addressing the first inquiry, Figure 7 compares how target communication interacts with the inflation coefficient of the Taylor rule, ϕ and its effects on the

aggregate loss. This figure shows the norm of the consumption wedges, which is defined by:

$$\left\|\mathcal{W}^R\right\|_\omega = \sqrt{\sum_g \omega_g \sum_t \left(\mathcal{W}^R_{g,t}\right)^2}. \tag{55}$$

and measures the cohort-weighted distance to the rational path after an inflationary shock. The HANK figures use the same norm with $\mathcal{W}^H$ instead of $\mathcal{W}^R$.

FIGURE 7. Target Communication and Traditional Monetary Policy



In the RANK model, panel A, a more aggressive rule lowers the loss, regardless of the communication intensity, since restrictive monetary policy reduces inflation and with it both the diagnostic and the credibility error in the same proportion. Communication removes the credibility error, but the offsetting part that is lost with it is worth about as much as the benefit at every ϕ, thus there is no substantial gain from fully intensive communication.

The story is different in the HANK model, as panel B shows. As in RANK, a more aggressive monetary policy rule reduces inflation on impact and lowers the loss, for a moderate value of ϕ. However, since the fall in labour income of the high-MPC households reduces aggregate demand even more, inflation falls faster and slightly below the target. Thus, cohorts that forecast with their own inflation persistence keep expecting inflation above the target even when it has returned to its original path. Contrary to RANK, the credibility error of the old cohorts adds to their diagnostic error instead of offsetting it, since experienced inflation is precisely what makes them expect higher inflation when it is actually decreasing. Thus, removing this over-prediction with intense communication is what makes a more aggressive rule worthwhile. The gain from full communication

rises from 3% of the loss with $\phi = 1.1$ to 7% with $\phi = 3$, so communication is worth more the more aggressive the rule is. Importantly, full target communication and a moderately aggressive ($\phi = 2.1$) monetary policy deliver the lowest welfare loss, which is evidence that both policies may be used as complements by the central bank.

Regarding the timing question, this paper considers the aggregate loss when full target communication ($\nu = 1$) starts k periods after the shock, against the case in which it starts on impact ($t = 0$). Let $\nu^{(k)}$ be the intensity path that adopts full target communication k periods after the shock hits the economy. The cost of the delay is

$$\Delta(k) \equiv \|\mathcal{W}^R\|_\omega(k) - \|\mathcal{W}^R\|_\omega(0), \tag{56}$$

where $\|\mathcal{W}^R\|_\omega(k)$ is the norm of (55) on the equilibrium path under $\nu^{(k)}$, so that $\Delta(k)$ is the extra loss of starting at k instead of on impact.

FIGURE 8. Timing of Target Communication

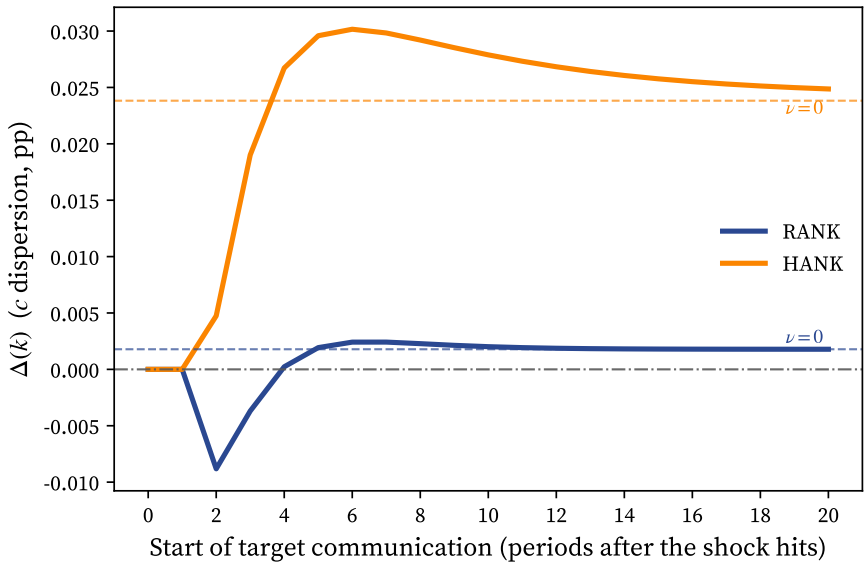


Figure 8 plots the cost of delayed communication and shows that in the RANK model a short delay is slightly better than communicating on impact. On impact every cohort under-reacts to the shock, and the credibility error of those with inflation experience above the target offsets part of that under-reaction. Keeping it for a short period, although it means keeping the whole credibility error, turns out to be beneficial. Once this period has passed, the older cohorts start to over-predict inflation and their credibility error grow together with the diagnostic error, so

waiting to communicate increases the cost. On the contrary, in the HANK model starting on impact is best. The offsetting on impact is also present there, but the over-prediction of the older cohorts right after it is stronger, so removing it early dominates. Thus, reacting early prevents credibility from eroding, so the target has to be communicated once the shock hits, not once its effects on expectations are visible.

6. Conclusion

Central banks respond to inflationary episodes by constantly repeating that they will do everything in their power to bring inflation back to the target. While this message seems to reach the most sophisticated individuals, households largely do not notice these announcements when forming their expectations. This paper asks what the central bank gains by communicating its target intensively. It builds a model of inflation expectations with endogenous credibility that is consistent with the under-reaction of households by incorporating DE as a fundamental part of how households process news. It then applies this mechanism to the New Keynesian model, both in its simple version and in its heterogeneous-agent version.

Three lessons emerge from this paper. The first is related to who pays in terms of welfare after an inflationary shock. In both models every cohort under-reacts to the shock, but the loss falls on the cohorts under 55, whose forecasts is barely affected. The old cohorts lose less because the persistent inflation they lived through makes their inflation expectations to increase after the shock, but the same history erodes their credibility and makes them over-predict inflation afterwards. The second lesson is that intensively communicating the target is not always optimal. Communication removes the credibility error, but with it the part of that error that was offsetting the under-reaction, so the old cohorts gain and the cohorts whose experience is only slightly above the target lose. The optimal intensity balances both effects and is higher in the HANK model, where the old cohorts over-estimate inflation for longer and the gain from anchoring them dominates.

The third lesson is about the connection between target-based communication and traditional monetary policy. With a representative agent, a more aggressive interest rate rule lowers the loss regardless of communication. On the other hand,

ith heterogeneous agents, the stronger demand contraction of constrained households brings inflation back to the target faster, the old cohorts keep over-predicting it, and communication is the tool that prevents it. Consequently, the target has to be communicated when the shock hits and not once its effects on credibility are visible, since communicating a year late is worse than not communicating at all.

This paper leaves some open questions. The communication intensity is a choice of the CB, but its reception by households is taken as given. Studying theoretically how the central bank can reach households with its communication is a research area that would complement this paper. Another pending area comes from relaxing the assumption that inflation experiences are fixed over time. Allowing recent inflation to be incorporated into experienced inflation presents a methodological challenge, but it would add realism to this analysis. Finally, extending the mechanism to the interaction between target communication and forward guidance on the policy rate would offer a wider spectrum of communication policies that the CB can use under distorted expectations.

References

Angeletos, G.-M. and Sastry, K. A. (2021). Managing expectations: Instruments versus targets. *Quarterly Journal of Economics*, 136(4):2467–2532.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–67.

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408.

Auclert, A., Rognlie, M., and Straub, L. (2024). The intertemporal keynesian cross. *Journal of Political Economy*, 132(12):4068–4121.

Auclert, A., Rognlie, M., and Straub, L. (2025). Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics*, 17:539–562.

Bank of England (2022). Monetary policy report, november 2022. Bank of England, 3 November 2022.

Bianchi, F., Ilut, C., and Saijo, H. (2024). Diagnostic business cycles. *Review of Economic Studies*, 91(1):129–162.

Binder, C. (2017). Fed speak on main street: Central bank communication and household expectations. *Journal of Macroeconomics*, 52:238–251.

Blinder, A. S. (2000). Central-bank credibility: Why do we care? how do we build it? *American Economic Review*, 90(5):1421–1431.

Bomfim, A. N. and Rudebusch, G. D. (2000). Opportunistic and deliberate disinflation under imperfect credibility. *Journal of Money, Credit and Banking*, 32(4):707–721.

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

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic expectations and credit cycles. *The Journal of Finance*, 73(1):199–227.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2022). Overreaction and diagnostic expectations in macroeconomics. *Journal of Economic Perspectives*, 36(3):223–244.

Bordo, M. D. and Siklos, P. L. (2015). Central bank credibility before and after the crisis. NBER Working Paper 21710, National Bureau of Economic Research, Cambridge, MA.

Born, B., Enders, Z., Menkhoff, M., Müller, G. J., and Niemann, K. (2025). Firm expectations and news: Micro v macro. Working paper, University of Bonn, Heidelberg University, University of Copenhagen and University of Tübingen.

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

Christelis, D., Georgarakos, D., Jappelli, T., and van Rooij, M. (2020). Trust in the central

bank and inflation expectations. *International Journal of Central Banking*, 16(6):1–37.

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

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

D'Acunto, F., Hoang, D., Paloviita, M., and Weber, M. (2020). Effective policy communication: Targets versus instruments. Becker Friedman Institute Working Paper 2020-148, University of Chicago.

Ehrmann, M., Georgarakos, D., and Kenny, G. (2025). Credibility gains from central bank communication with the public. *European Economic Review*, 177:105069.

Francis, N. (2026). The perceived inflation wedge. NBER Working Paper 35354, National Bureau of Economic Research.

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review*, 110(8):2271–2327.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press, 2 edition.

Gennaioli, N., Leva, M., Schoenle, R., and Shleifer, A. (2024). How inflation expectations de-anchor: The role of selective memory cues. NBER Working Paper 32633, National Bureau of Economic Research, Cambridge, MA.

Goy, G., Hommes, C., and Mavromatis, K. (2022). Forward guidance and the role of central bank credibility under heterogeneous beliefs. *Journal of Economic Behavior and Organization*, 200:1240–1274.

Guerreiro, J. (2023). Belief disagreement and business cycles. Working paper, Northwestern University.

Haldane, A. and McMahon, M. (2018). Central bank communications and the general public. *AEA Papers and Proceedings*, 108:578–583.

Hommes, C. and Lustenhouwer, J. (2019). Inflation targeting and liquidity traps under endogenous credibility. *Journal of Monetary Economics*, 107:48–62.

Koop, G. and Korobilis, D. (2012). Forecasting inflation using dynamic model averaging. *International Economic Review*, 53(3):867–886.

Kumar, S., Afrouzi, H., Coibion, O., and Gorodnichenko, Y. (2015). Inflation targeting does not anchor inflation expectations: Evidence from firms in new zealand. *Brookings Papers on Economic Activity*, 2015(2):151–225.

Lagarde, C. (2022). Monetary policy in a high inflation environment: Commitment and clarity. Lecture organised by Eesti Pank, Tallinn, 4 November 2022.

Lamla, M. J. and Vinogradov, D. V. (2019). Central bank announcements: Big news for little people? *Journal of Monetary Economics*, 108:21–38.

L'Huillier, J.-P., Singh, S. R., and Yoo, D. (2024). Incorporating diagnostic expectations into the new keynesian framework. *Review of Economic Studies*, 91(5):3013–3046.

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

McMahon, M. and Rholes, R. (2025). Building central bank credibility: The role of forecast performance. Working paper.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.

Powell, J. H. (2022). Monetary policy and price stability. Speech at the Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City, 26 August 2022.

Powell, J. H. (2023). Inflation: Progress and the path ahead. Speech at the Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City, 25 August 2023.

Romer, C. D. and Romer, D. H. (2000). Federal reserve information and the behavior of interest rates. *American Economic Review*, 90(3):429–457.

Salle, I., Gorodnichenko, Y., and Coibion, O. (2023). Lifetime memories of inflation: Evidence from surveys and the lab. NBER Working Paper 31996, National Bureau of Economic Research, Cambridge, MA.

Shabalina, E. and Tzaawa-Krenzler, M. (2024). Heterogeneous attention to inflation and monetary policy. Rba workshop paper, Reserve Bank of Australia.

Young, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell-smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control*, 34(1):36–41.

Appendix A. Proofs

A.1. Proposition 1

PROOF. The first step is to write the forecast error of each hypothesis. Under Equation (3) with anchor m_h, the forecast of π_t made at $t-1$ is $\rho_g \pi_{t-1} + (1-\rho_g) m_h$, so the error is Equation (5),

$$e_{h,t} = \pi_t - \rho_g \pi_{t-1} - (1-\rho_g)\, m_h, \qquad h \in \{T, E\}.$$

By Assumption 6, these are the errors the cohort uses to compare the two hypotheses. The next step is to write the discounted log-likelihood of each hypothesis in terms of these errors. Under Equation (3), $\pi_t \mid \pi_{t-1}, h \sim N(\rho_g \pi_{t-1} + (1-\rho_g) m_h, \sigma_g^2)$, whose log-density is Equation (6), $\log f(\pi_t \mid \pi_{t-1}, h) = -\frac{1}{2}\log(2\pi\sigma_g^2) - e_{h,t}^2/(2\sigma_g^2)$. Substituting into Equation (7) and using $\sum_{j\geq 0} \rho_U^j = 1/(1-\rho_U)$,

$$\ell_{h,t} = c - \frac{1}{2\sigma_g^2} \sum_{j\geq 0} \rho_U^j\, e_{h,t-j}^2, \qquad c \equiv -\frac{\log(2\pi\sigma_g^2)}{2(1-\rho_U)}. \tag{A1}$$

The constant c is common across hypotheses. Next, define the performance of hypothesis h as the discounted average of its squared errors,

$$U_{h,t} \equiv -(1-\rho_U) \sum_{j\geq 0} \rho_U^j\, e_{h,t-j}^2, \tag{A2}$$

Manipulating the sum,

$$U_{h,t} = -(1-\rho_U)\, e_{h,t}^2 + \rho_U \Big[-(1-\rho_U) \sum_{j\geq 1} \rho_U^{j-1}\, e_{h,t-j}^2 \Big] = -(1-\rho_U)\, e_{h,t}^2 + \rho_U\, U_{h,t-1},$$

which is Equation (9). Comparing (A1) with (A2), the sum in (A1) equals $-U_{h,t}/(1-\rho_U)$, so

$$\ell_{h,t} = c + \frac{U_{h,t}}{2\sigma_g^2(1-\rho_U)} = c + b\, U_{h,t}, \qquad b \equiv \frac{1}{2\sigma_g^2(1-\rho_U)}. \tag{A3}$$

Finally, I apply Bayes' rule to Equation A3 to get the updating. By Assumption 4 the prior is 1/2 on each hypothesis and the likelihood of hypothesis h is $e^{\ell_{h,t-1}}$, with

the information up to $t-1$, which leads the result,

$$\chi_{g,t} = \frac{\frac{1}{2}e^{\ell_{T,t-1}}}{\frac{1}{2}e^{\ell_{T,t-1}} + \frac{1}{2}e^{\ell_{E,t-1}}} = \left[1 + e^{-(\ell_{T,t-1} - \ell_{E,t-1})}\right]^{-1}.$$

□

A.2. Lemma 1

PROOF. The first step is the predictive distribution under each hypothesis. Under Equation (3) with anchor m_h,

$$\pi_{t+1} \mid I_t, h \;\sim\; N(\mu_{h,t}, \sigma_g^2), \qquad \mu_{h,t} = \rho_g \pi_t + (1-\rho_g)\, m_h,$$

because π_t is known at t and $\varepsilon_{t+1} \sim N(0, \sigma_g^2)$. By Assumption 2 the cohort does not know the anchor, and by Assumption 3 it considers the two values m_T and $m_{E,g}$. By Proposition 1 it attaches probability $\chi_{g,t}$ to T and $1-\chi_{g,t}$ to E. Therefore

$$f(\pi_{t+1} \mid I_t) = \chi_{g,t}\, f(\pi_{t+1} \mid I_t, T) + (1-\chi_{g,t})\, f(\pi_{t+1} \mid I_t, E),$$

which is the mix ot the two equations in (10). Further, calculate its mean as the weighted sum of the two conditional means,

$$\mathbb{E}_{g,t}[\pi_{t+1}] = \chi_{g,t}\, \mu_{T,t} + (1-\chi_{g,t})\, \mu_{E,t} = \rho_g\, \pi_t + (1-\rho_g)\big[\chi_{g,t}\, m_T + (1-\chi_{g,t})\, m_{E,g}\big].$$

Finally, substituting $m_T = \pi^*$ and $m_{E,g} = \nu\pi^* + (1-\nu)\bar{\pi}_g$ from Equation (4),

$$\begin{aligned} m_{g,t} &= \chi_{g,t}\, \pi^* + (1-\chi_{g,t})\big[\nu\pi^* + (1-\nu)\bar{\pi}_g\big] \\ &= \pi^* + (1-\chi_{g,t})(1-\nu)\big(\bar{\pi}_g - \pi^*\big) = \pi^* + (1-\chi_{g,t})(1-\nu)\,\Delta_g, \end{aligned}$$

which is what we wanted to prove. □

A.3. Proposition 2

PROOF. The proof has four steps. The first step finds the steady state of credibility. The second step linearises the law of credibility around it. The third step writes the Bayesian forecast as the first row of a linear state. The fourth step applies the diagnostic operator to that forecast. Throughout, a hat denotes a deviation

from the steady state $\pi_t = \pi^*$, and we write $u_t \equiv \pi_t - \rho_g \pi_{t-1}$, so that the error of Equation (5) reads $e_{h,t} = u_t - (1-\rho_g) m_h$.

Step 1: steady state of credibility. At $\pi_t = \pi_{t-1} = \pi^*$ we have $u = (1-\rho_g)\pi^*$, and Equation (5) with the anchors of Equation (4) gives,

$$e_T = (1-\rho_g)(\pi^* - m_T) = 0, \qquad e_E = (1-\rho_g)(\pi^* - m_{E,g}) = -(1-\rho_g)(1-\nu)\,\Delta_g,$$

using $\pi^* - m_{E,g} = (1-\nu)(\pi^* - \bar{\pi}_g) = -(1-\nu)\Delta_g$. Thus, the steady state value of the recursion is $U_h = -e_h^2$. Therefore

$$\Delta U \equiv U_T - U_E = e_E^2 - e_T^2 = (1-\rho_g)^2 (1-\nu)^2 \,\Delta_g^2 \;\geq\; 0. \tag{A4}$$

Hence, substituting (A4) into Equation (8), the steady-state credibility is

$$\chi_g \;=\; \left[1 + e^{-b\,(1-\rho_g)^2 (1-\nu)^2\,\Delta_g^2}\right]^{-1} \;\geq\; \tfrac{1}{2}. \tag{A5}$$

Step 2: linearised law of credibility. Subtracting the two recursions of the different hypothesis in (9),

$$\Delta U_t = (1-\rho_U)\left(e_{E,t}^2 - e_{T,t}^2\right) + \rho_U\,\Delta U_{t-1}. \tag{A6}$$

we can see it is constant,

$$e_{E,t} - e_{T,t} = (1-\rho_g)(m_T - m_{E,g}) = -(1-\rho_g)(1-\nu)\,\Delta_g, \tag{A7}$$

and their sum is $e_{E,t} + e_{T,t} = 2u_t - (1-\rho_g)(m_T + m_{E,g})$. Therefore, multiplying the two,

$$e_{E,t}^2 - e_{T,t}^2 = -(1-\rho_g)(1-\nu)\Delta_g \left[2u_t - (1-\rho_g)(m_T + m_{E,g})\right], \tag{A8}$$

which at the steady state returns (A4). Substituting (A8) into (A6) and differentiating around the steady state,

$$d\Delta U_t = -2(1-\rho_U)(1-\rho_g)(1-\nu)\Delta_g\, du_t + \rho_U\, d\Delta U_{t-1}, \qquad du_t = (\hat{\pi}_t - \rho_g \hat{\pi}_{t-1}), \tag{A9}$$

Next, I differentiate Equation (8). Observe the logistic function $\Lambda(x) = [1+e^{-x}]^{-1}$ has derivative $\Lambda(1-\Lambda)$, so

$$\hat{\chi}_{g,t} = b\,\chi_g(1-\chi_g)\, d\Delta U_{t-1}, \tag{A10}$$

with χ_g the steady-state credibility of (A5). Therefore,

$$\begin{aligned}\hat{\chi}_{g,t} - \rho_U \hat{\chi}_{g,t-1} &= b\chi_g(1-\chi_g)\big[d\Delta U_{t-1} - \rho_U\, d\Delta U_{t-2}\big] \\ &= -2b\chi_g(1-\chi_g)(1-\rho_U)(1-\rho_g)(1-\nu)\Delta_g\, du_{t-1}.\end{aligned}$$

and substituting $du_{t-1} = (\hat{\pi}_{t-1} - \rho_g \hat{\pi}_{t-2})$ yields the linearised law of credibility,

$$\begin{aligned}\hat{\chi}_{g,t} &= \rho_U\, \hat{\chi}_{g,t-1} - \xi_g\big(\hat{\pi}_{t-1} - \rho_g\, \hat{\pi}_{t-2}\big), \\ \xi_g &= 2b\chi_g(1-\chi_g)\,(1-\rho_U)\,(1-\rho_g)\,(1-\nu)\,\Delta_g \times 100.\end{aligned} \tag{A11}$$

Step 3: the Bayesian forecast as a linear state. Observe $\hat{m}_{g,t} = \kappa_g\, \hat{\chi}_{g,t}$, with κ_g the credibility sensitivity of Equation (15). Therefore, the Bayesian forecast of Equation (11) in deviations is

$$\mathbb{E}_{g,t}[\hat{\pi}_{t+1}] = \rho_g\, \hat{\pi}_t + (1-\rho_g)\, \kappa_g\, \hat{\chi}_{g,t}. \tag{A12}$$

The cohort perceives two laws of motion. The first is the law of inflation (3) with the anchor at its posterior mean, which in deviations is $\hat{\pi}_{t+1} = \rho_g \hat{\pi}_t + (1-\rho_g)\kappa_g \hat{\chi}_{g,t} + \hat{\varepsilon}_{t+1}$. The second is the law of credibility (A11) at $t+1$, $\hat{\chi}_{g,t+1} = \rho_U \hat{\chi}_{g,t} - \xi_g \hat{\pi}_t + \xi_g \rho_g \hat{\pi}_{t-1}$. Stacking $z_{g,t} = (\hat{\pi}_t, \hat{\pi}_{t-1}, \hat{\chi}_{g,t})'$, the two laws are the perceived transition

$$z_{g,t+1} = M_g\, z_{g,t} + \hat{\varepsilon}_{t+1}\, e_1, \qquad M_g = \begin{pmatrix} \rho_g & 0 & (1-\rho_g)\kappa_g \\ 1 & 0 & 0 \\ -\xi_g & \xi_g \rho_g & \rho_U \end{pmatrix}, \tag{A13}$$

with e_1 the first unit vector. The first row of $M_g z_{g,t}$ is (A12), so $\mathbb{E}_{g,t}[\hat{\pi}_{t+1}] = e_1' M_g z_{g,t}$. Since (A13) is linear and $\hat{\varepsilon}_{t+1}$ is normal, the predictive distribution of $\hat{\pi}_{t+1}$ given I_t is normal with mean $e_1' M_g z_{g,t}$ and the variance of $\hat{\varepsilon}_{t+1}$.

Step 4: the diagnostic operator. By Assumption 5 the cohort reports the diagnostic version (12) of this predictive distribution, and by Assumption 6 the credibility it carries is the Bayesian $\hat{\chi}_{g,t}$ of (A11), not a distorted one. We first apply the operator to two normal densities with the same variance. Let $f(\cdot \mid I_t)$ be normal with mean μ_t and $f(\cdot \mid I_{t-1})$ be normal with mean μ_{t-1}, both with variance σ^2. Taking logs of (12), the terms that depend on x are

$$-\frac{1}{2\sigma^2}\Big[(1+\theta)(x-\mu_t)^2 - \theta(x-\mu_{t-1})^2\Big].$$

Expanding the squares, the coefficient on x^2 is $(1+\theta)-\theta=1$ and the coefficient on x is $-2[(1+\theta)\mu_t-\theta\mu_{t-1}]$. Hence, with $m\equiv(1+\theta)\mu_t-\theta\mu_{t-1}$,

$$\begin{aligned}(1+\theta)(x-\mu_t)^2-\theta(x-\mu_{t-1})^2 &= (x-m)^2+\left[(1+\theta)\mu_t^2-\theta\mu_{t-1}^2-m^2\right]\\ &= (x-m)^2-\theta(1+\theta)(\mu_t-\mu_{t-1})^2.\end{aligned}$$

The last term does not involve x, so the normalisation absorbs it. Therefore the diagnostic distribution is normal with the same variance σ^2 and mean

$$m=\mu_t+\theta\,(\mu_t-\mu_{t-1}),$$

which is Equation (13).

Next, I identify the two means. The current mean is $\mu_t=\mathbb{E}_{g,t}[\hat{\pi}_{t+1}]=e_1'M_g z_{g,t}$ from Step 3. Following Bordalo et al. (2018), the reference density $f(\pi_{t+1}\mid I_{t-1})$ is the one-step-ahead density of the state evaluated at the state the cohort expected at $t-1$, $\mathbb{E}_{g,t-1}[z_{g,t}]=M_g z_{g,t-1}$. It has the same variance as the current density and mean $\mu_{t-1}=e_1'M_g\,(M_g z_{g,t-1})=e_1'M_g^2 z_{g,t-1}$. Substituting both means into m,

$$\mathbb{E}^{\theta}_{g,t}[\hat{\pi}_{t+1}]=e_1'M_g z_{g,t}+\theta\big(e_1'M_g z_{g,t}-e_1'M_g^2 z_{g,t-1}\big). \tag{A14}$$

It remains to compute $e_1'M_g^2 z_{g,t-1}$. Observe,

$$M_g z_{g,t-1}=\begin{pmatrix}\rho_g\hat{\pi}_{t-1}+(1-\rho_g)\kappa_g\hat{\chi}_{g,t-1}\\ \hat{\pi}_{t-1}\\ \rho_U\hat{\chi}_{g,t-1}-\xi_g\hat{\pi}_{t-1}+\xi_g\rho_g\hat{\pi}_{t-2}\end{pmatrix}=\begin{pmatrix}\mathbb{E}_{g,t-1}[\hat{\pi}_t]\\ \hat{\pi}_{t-1}\\ \hat{\chi}_{g,t}\end{pmatrix}.$$

Multiplying by the first row of M_g,

$$e_1'M_g^2 z_{g,t-1}=\rho_g\,\mathbb{E}_{g,t-1}[\hat{\pi}_t]+(1-\rho_g)\kappa_g\,\hat{\chi}_{g,t}.$$

Subtracting this from $e_1'M_g z_{g,t}$ in (A12), the credibility terms cancel because both carry $\hat{\chi}_{g,t}$,

$$e_1'M_g z_{g,t}-e_1'M_g^2 z_{g,t-1}=\rho_g\big(\hat{\pi}_t-\mathbb{E}_{g,t-1}[\hat{\pi}_t]\big)=\rho_g\big(\hat{\pi}_t-\rho_g\hat{\pi}_{t-1}-(1-\rho_g)\kappa_g\hat{\chi}_{g,t-1}\big),$$

using (A12) dated $t-1$ once more. Substituting this difference and (A12) into (A14) yields Equation (14), which is what we wanted to prove.

Finally, the same expression holds in levels. A direct application of the diagnostic operator leads to

$$\mathbb{E}^{\theta}_{g,t}[\pi_{t+1}] \;=\; \rho_g\,\pi_t + (1-\rho_g)\,m_{g,t} + \theta\,\rho_g\big[u_t - (1-\rho_g)\,m_{g,t-1}\big], \tag{A15}$$

where

$$\begin{aligned} u_t - (1-\rho_g)\,m_{g,t-1} &= u_t - (1-\rho_g)\big[\chi_{g,t-1}\,m_T + (1-\chi_{g,t-1})\,m_{E,g}\big] \\ &= \chi_{g,t-1}\,e_{T,t} + (1-\chi_{g,t-1})\,e_{E,t}. \end{aligned}$$

□

A.4. Lemma 2

PROOF. The first step is to log-linearise the cohort-g Euler equation. In levels it is

$$C_{g,t}^{-1/\sigma} \;=\; \beta_g(1+r^{\text{per}}_{g,t})\,\mathbb{E}^{\theta}_{g,t}[C_{g,t+1}^{-1/\sigma}].$$

where uppercase letters denote levels and $r^{\text{per}}_{g,t}$ is the perceived rate of (18). Following the convention of the paper, lowercase letters denote first-order deviations from the steady state $\{C_g^{ss},\, r_g^{per\,ss} = \beta_g^{-1} - 1,\, a_g^{ss} = 0,\, \pi^{ss} = \pi^*,\, \chi_g^{ss}\}$, so $c_{g,t} \equiv \log(C_{g,t}/C_g^{ss})$, $\hat{r}^{\text{per}}_{g,t} \equiv r^{\text{per}}_{g,t} - r_g^{per\,ss}$, and analogously for y and a. Taking logs of the Euler equation dated $t+h$ and keeping first-order terms,

$$c_{g,t+h+1} - c_{g,t+h} = \sigma\,\hat{r}^{\text{per}}_{g,t+h}. \tag{A16}$$

Forward iterating (A16) from t to $t+h$ yields

$$c_{g,t+h} = c_{g,t} + \sigma \sum_{k=0}^{h-1} \hat{r}^{\text{per}}_{g,t+k}. \tag{A17}$$

The next step is to find the log-linear version of the intertemporal budget constraint. Log-linearising (17) around the same steady state, with $1 + r^{ss} = \beta_g^{-1}$ and $a_{g,t}$ measured in units of steady-state consumption,

$$a_{g,t+1} = \beta_g^{-1} a_{g,t} + y_t - c_{g,t}. \tag{A18}$$

Multiplying (A18) at $t+h$ by β_g^{h+1} and summing from $h=0$ to $h=T-1$,

$$\sum_{h=0}^{T-1} \beta_g^{h+1} a_{g,t+h+1} = \sum_{h=0}^{T-1} \beta_g^{h} a_{g,t+h} + \sum_{h=0}^{T-1} \beta_g^{h+1} \big(y_{t+h} - c_{g,t+h} \big).$$

The two sums in a share every term except the last one on the left and the first one on the right, so after cancelling the common terms

$$\beta_g^{T} a_{g,t+T} - a_{g,t} = \sum_{h=0}^{T-1} \beta_g^{h+1} \big(y_{t+h} - c_{g,t+h} \big).$$

Taking $T \to \infty$, imposing the transversality condition $\lim_{T\to\infty} \beta_g^{T} a_{g,t+T} = 0$, and recalling that the cohort starts at the steady state, $a_{g,t} = 0$, we reach the familiar expression

$$\sum_{h=0}^{\infty} \beta_g^{h} c_{g,t+h} = \sum_{h=0}^{\infty} \beta_g^{h} y_{t+h}. \tag{A19}$$

The next step is to substitute the consumption path (A17) into (A19). Using $\sum_{h\geq 0} \beta_g^{h} = 1/(1-\beta_g)$ for the term in $c_{g,t}$,

$$\frac{1}{1-\beta_g} c_{g,t} = \sum_{h=0}^{\infty} \beta_g^{h} y_{t+h} - \sigma \sum_{h=0}^{\infty} \beta_g^{h} \sum_{k=0}^{h-1} \hat{r}^{\text{per}}_{g,t+k}. \tag{A20}$$

Now observe that the double sum collects, for each date $t+k$, the discount factors of all the later dates $h \geq k+1$,

$$\begin{aligned} \sum_{h=0}^{\infty} \beta_g^{h} \sum_{k=0}^{h-1} \hat{r}^{\text{per}}_{g,t+k} &= \hat{r}^{\text{per}}_{g,t} \big(\beta_g + \beta_g^2 + \beta_g^3 + \cdots \big) + \hat{r}^{\text{per}}_{g,t+1} \big(\beta_g^2 + \beta_g^3 + \cdots \big) \\ &\quad + \hat{r}^{\text{per}}_{g,t+2} \big(\beta_g^3 + \beta_g^4 + \cdots \big) + \cdots \\ &= \sum_{k=0}^{\infty} \hat{r}^{\text{per}}_{g,t+k} \sum_{h=k+1}^{\infty} \beta_g^{h} = \sum_{k=0}^{\infty} \frac{\beta_g^{k+1}}{1-\beta_g} \hat{r}^{\text{per}}_{g,t+k}. \end{aligned}$$

Substituting into (A20),

$$\frac{1}{1-\beta_g} c_{g,t} = \sum_{h=0}^{\infty} \beta_g^{h} y_{t+h} - \sigma \sum_{h=0}^{\infty} \frac{\beta_g^{h+1}}{1-\beta_g} \hat{r}^{\text{per}}_{g,t+h}.$$

Finally, multiplying by $1-\beta_g$ proves the result

$$c_{g,t} \;=\; (1-\beta_g)\sum_{h=0}^{\infty}\beta_g^h\, y_{t+h} \;-\; \sigma\sum_{h=0}^{\infty}\beta_g^{h+1}\,\hat{r}^{\text{per}}_{g,t+h},$$

□

A.5. Proposition 3

PROOF. The first step is to write the cohort-g consumption path. By Lemma 2, the log-linearised Euler equation is $dc_{g,t} = dc_{g,t+1} - \sigma\,\hat{r}^{\text{per}}_{g,t}$, where, $\hat{r}^{\text{per}}_{g,t} = \hat{\imath}_t - \mathrm{d}\mathbb{E}^{\theta}_{g,t}[\hat{\pi}_{t+1}]$. By Proposition 2 $\mathrm{d}\mathbb{E}^{\theta}_{g,t}[\hat{\pi}_{t+1}] = \Lambda^{\pi}_g\hat{\pi}_t + \Lambda^{\chi}_g\hat{\chi}_{g,t}$, so

$$dc_{g,t} = dc_{g,t+1} - \sigma\big(\hat{\imath}_t - \Lambda^{\pi}_g\hat{\pi}_t - \Lambda^{\chi}_g\hat{\chi}_{g,t}\big).$$

Forward iterating and using $\lim_{s\to\infty} dc_{g,t+s} = 0$ yields

$$dc_{g,t} = -\sigma\sum_{s\geq 0}\big(\hat{\imath}_{t+s} - \Lambda^{\pi}_g\hat{\pi}_{t+s} - \Lambda^{\chi}_g\hat{\chi}_{g,t+s}\big). \tag{A21}$$

The next step is to get the rational path. It solves the same Euler equation on the same paths of the nominal rate and inflation, with the *realised inflation* instead of the forecast, $dc^{rat}_t = dc^{rat}_{t+1} - \sigma(\hat{\imath}_t - \hat{\pi}_{t+1})$. This equation is cohort independent, so forward iterating with the same terminal condition gives, for every g,

$$dc^{rat}_t = -\sigma\sum_{s\geq 0}\big(\hat{\imath}_{t+s} - \hat{\pi}_{t+s+1}\big).$$

Next, I calculate the consumption wedge. Subtracting the rational path from (A21), the nominal rate cancels and

$$\mathcal{W}^R_{g,t} = \sigma\sum_{s\geq 0}\big(\Lambda^{\pi}_g\hat{\pi}_{t+s} + \Lambda^{\chi}_g\hat{\chi}_{g,t+s} - \hat{\pi}_{t+s+1}\big) = \sigma\sum_{s\geq 0}\Big(\mathrm{d}\mathbb{E}^{\theta}_{g,t+s}[\hat{\pi}_{t+s+1}] - \hat{\pi}_{t+s+1}\Big),$$

the sum of the cohort's one-step-ahead forecast errors. Grouping the terms in $\hat{\pi}$ and the terms in $\hat{\chi}_g$ gives (42). The first group, $\Lambda^{\pi}_g\hat{\pi}_{t+s} - \hat{\pi}_{t+s+1}$, is the error the cohort would make with the credibility channel off, $\kappa_g = 0$. The second group, $\Lambda^{\chi}_g\hat{\chi}_{g,t+s}$, is what credibility adds to the forecast.

The next step is to write the wedge in terms of the two errors of (43). The first group, summed over s and multiplied by σ, is $\mathcal{A}_{g,t}$ by definition. For the second

group, by (41), $\Lambda^{\chi}_{g}\hat{\chi}_{g,t} = (1-\rho_g)\kappa_g X_{g,t}$ with $X_{g,t} \equiv \hat{\chi}_{g,t} - \theta\rho_g\hat{\chi}_{g,t-1}$, and by (15), $\kappa_g = -(1-\nu)\Delta_g$. Substituting,

$$\sigma \sum_{s\geq 0} \Lambda^{\chi}_{g}\hat{\chi}_{g,t+s} = (1-\nu)\Big[-\sigma \sum_{s\geq 0}(1-\rho_g)\,\Delta_g\, X_{g,t+s}\Big] = (1-\nu)\,\mathcal{B}_{g,t}.$$

Therefore $\mathcal{W}^{R}_{g,t} = \mathcal{A}_{g,t} + (1-\nu)\mathcal{B}_{g,t}$. Note that $\mathcal{A}_{g,t}$ does not depend on ν, because Λ^{π}_{g} depends only on (ρ_g, θ) and ν moves neither of them (Assumption 2).

The next step is the welfare loss. Take cohort g's discounted utility along a consumption path $\{c_{g,t+s}\}_{s\geq 0}$,

$$U_g\big(\{c_{g,t+s}\}\big) = \sum_{s\geq 0} \beta_g^s\, u(c_{g,t+s}), \tag{A22}$$

and expand it to second order around the rational path,

$$\begin{aligned} U_g\big(\{c_{g,t+s}\}\big) = U_g\big(\{c^{rat}_{g,t+s}\}\big) &+ \sum_{s\geq 0} \beta_g^s\, u'(c^{rat}_{g,t+s})\big(c_{g,t+s} - c^{rat}_{g,t+s}\big) \\ &+ \tfrac{1}{2}\sum_{s\geq 0} \beta_g^s\, u''(c^{rat}_{g,t+s})\big(c_{g,t+s} - c^{rat}_{g,t+s}\big)^2 + O\big(\|c - c^{rat}\|^3\big). \end{aligned}$$

Let u'' be evaluated at the steady state, so $-u''(c^{rat}_{g,t+s}) = A_g > 0$ at every date. Therefore,

$$U_g\big(\{c^{rat}_{g,t+s}\}\big) - U_g\big(\{c_{g,t+s}\}\big) = A_g \sum_{s\geq 0} \beta_g^s\, \tfrac{1}{2}\big(\mathcal{W}^{R}_{g,t+s}\big)^2.$$

which leads

$$\frac{U_g\big(c^{rat}_{g,t}\big) - U_g\big(c_{g,t}\big)}{A_g} = \tfrac{1}{2}\big(\mathcal{W}^{R}_{g,t}\big)^2 = L^{R}_{g,t}.$$

Substituting $\mathcal{W}^{R}_{g,t} = \mathcal{A}_{g,t} + (1-\nu)\mathcal{B}_{g,t}$ and expanding the square,

$$\tfrac{1}{2}\big(\mathcal{W}^{R}_{g,t}\big)^2 = \tfrac{1}{2}\mathcal{A}^2_{g,t} + (1-\nu)\,\mathcal{A}_{g,t}\mathcal{B}_{g,t} + \tfrac{1}{2}(1-\nu)^2\mathcal{B}^2_{g,t},$$

which is the second equality of (44). □

A.6. Proposition 4

PROOF. The first step is to write the cohort-g consumption path. Recall (35),

$$\mathrm{d}\mathbf{C}_g = \mathcal{J}^{C,r}_{g}\mathrm{d}\mathbf{r} + \mathcal{J}^{C,r^{\mathrm{per}}}_{g}\mathrm{d}\mathbf{r}^{\mathrm{per}}_{g} + \mathcal{J}^{C,Y}_{g}\mathrm{d}\mathbf{Y}_g + \mathcal{J}^{C,\tau}_{g}\mathrm{d}\boldsymbol{\tau}.$$

By (36), the perceived-rate term splits into a Fisher part, $\mathcal{J}_g^{C,r^{\text{per}}}\mathrm{d}\mathbf{i}$, and a belief part, $-\mathcal{J}_g^{C,r^{\text{per}}}\left(\Lambda_g^{\pi}\mathrm{d}\boldsymbol{\pi}+\Lambda_g^{\chi}\mathrm{d}\boldsymbol{\chi}_g\right)=-\mathcal{J}_g^{C,r^{\text{per}}}\mathrm{d}\mathcal{E}_g$, where $\Lambda_g^{\pi}=\rho_g[(1+\theta)\mathbb{I}-\theta\rho_g L]$ and $\Lambda_g^{\chi}=(1-\rho_g)\kappa_g[\mathbb{I}-\theta\rho_g L]$ are the operators of (41) stacked over horizons, and $\mathrm{d}\mathcal{E}_g$ is the path of the cohort's forecast. Dividing by C_g, writing $\mathcal{J}_g^{c,X}\equiv\mathcal{J}_g^{C,X}/C_g$ and $\mathcal{J}_g^{c,\pi}\equiv-\mathcal{J}_g^{c,r^{\text{per}}}$, and putting the Fisher part together with income and taxes, the per-capita path is

$$\mathrm{d}\mathbf{c}_g=\mathcal{J}_g^{c,\pi}\,\mathrm{d}\mathcal{E}_g+\mathcal{J}_g^{c,r}\,\mathrm{d}\mathbf{r}+\sum_{X\in\{i,Y,\tau\}}\mathcal{J}_g^{c,X}\,\mathrm{d}\mathbf{X}. \tag{A23}$$

The next step is the rational path of cohort g. It is the same path with one change: the subjective forecast $\mathrm{d}\mathcal{E}_g$ is replaced by the realised next-period inflation, $\mathbb{F}\mathrm{d}\boldsymbol{\pi}$. The paths of the real rate, the nominal rate, income and taxes are the same. The Jacobians are also the same, because they are evaluated at the steady state, where the perceived and the realised rate coincide (Appendix B.3). Therefore

$$\mathrm{d}\mathbf{c}_g^{rat}=\mathcal{J}_g^{c,\pi}\,\mathbb{F}\,\mathrm{d}\boldsymbol{\pi}+\mathcal{J}_g^{c,r}\,\mathrm{d}\mathbf{r}+\sum_{X\in\{i,Y,\tau\}}\mathcal{J}_g^{c,X}\,\mathrm{d}\mathbf{X}. \tag{A24}$$

To get the consumption wedge, subtract (A24) from (A23). Observe the real-rate term and the $\{i,Y,\tau\}$ terms cancel, and only the forecast term remains,

$$\mathcal{W}_{g,t}^{H}=\left[\mathcal{J}_g^{c,\pi}\left(\mathrm{d}\mathcal{E}_g-\mathbb{F}\mathrm{d}\boldsymbol{\pi}\right)\right]_t.$$

The date-t element of $\mathrm{d}\mathcal{E}_g-\mathbb{F}\mathrm{d}\boldsymbol{\pi}$ is $\mathrm{d}\mathbb{E}_{g,t}^{\theta}[\pi_{t+1}]-\mathrm{d}\pi_{t+1}$, the forecast error of the cohort, so the wedge is the forecast error scaled by the cohort's own consumption response. Substituting $\mathrm{d}\mathcal{E}_g=\Lambda_g^{\pi}\mathrm{d}\boldsymbol{\pi}+\Lambda_g^{\chi}\mathrm{d}\boldsymbol{\chi}_g$ and grouping the terms in $\mathrm{d}\boldsymbol{\pi}$ and in $\mathrm{d}\boldsymbol{\chi}_g$ gives (46). As in RANK, the first group is the error with the credibility channel off, $\kappa_g=0$ while the second group is what credibility adds. When every cohort has the same per-capita Jacobian and it equals the permanent-income response, $[\mathcal{J}_g^{c,\pi}\mathbf{x}]_t=\sigma\sum_{s\geq 0}x_{t+s}$, (46) is the RANK wedge (42) and the rational path is the same for every cohort.

The next step is to write the wedge in terms of the two errors of (47) as in the RANK model. The first group of (46) is $\mathcal{A}_{g,t}^{H}$ by definition. For the second group, $\Lambda_g^{\chi}\mathrm{d}\boldsymbol{\chi}_g=(1-\rho_g)\kappa_g\mathbf{X}_g$ with $\mathbf{X}_g\equiv(\mathbb{I}-\theta\rho_g L)\mathrm{d}\boldsymbol{\chi}_g$, and by (15), $\kappa_g=-(1-\nu)\Delta_g$.

Substituting,

$$\left[\mathcal{J}_g^{c,\pi}\Lambda_g^{\chi}\mathrm{d}\boldsymbol{\chi}_g\right]_t = (1-\nu)\Big[-(1-\rho_g)\,\Delta_g\left[\mathcal{J}_g^{c,\pi}\mathbf{X}_g\right]_t\Big] = (1-\nu)\,\mathcal{B}_{g,t}^{H},$$

so $\mathcal{W}_{g,t}^{H} = \mathcal{A}_{g,t}^{H} + (1-\nu)\mathcal{B}_{g,t}^{H}$. Similarly, $\mathcal{A}_{g,t}^{H}$ does not depend on ν. Note that the credibility path $\mathrm{d}\boldsymbol{\chi}_g$ is cohort-specific, so the credibility error differs across cohorts even when the Jacobians are the same for all of them.

Next, I turn into the welfare loss. Expanding cohort g's discounted utility to second order around its own rational path, as in the proof of Proposition 3, we get the cohort loss of (40). Substituting $\mathcal{W}_{g,t}^{H} = \mathcal{A}_{g,t}^{H} + (1-\nu)\mathcal{B}_{g,t}^{H}$ and expanding the square gives (48). Finally, multiplying (48) by ω_g and summing across cohorts gives $\mathcal{L}_t^{H} = \sum_g \omega_g L_{g,t}^{H}$ as in (49), which is what we wanted to prove. □

A.7. Proposition 5

PROOF. *Part 1 (RANK).* By Proposition 3, $\mathcal{W}_{g,t}^{R}(\nu) = \mathcal{A}_{g,t} + (1-\nu)\mathcal{B}_{g,t}$, with $\mathcal{A}_{g,t}$ and $\mathcal{B}_{g,t}$ as in (43). The realised paths are given, so $\mathcal{A}_{g,t}$ and $\mathcal{B}_{g,t}$ are constants in the problem, and neither depends on ν. Summing the cohort loss (44) over t,

$$L_g^{R}(\nu) = \tfrac{1}{2}\sum_{t\geq 0}\mathcal{A}_{g,t}^{2} + (1-\nu)\sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t} + \tfrac{1}{2}(1-\nu)^{2}\sum_{t\geq 0}\mathcal{B}_{g,t}^{2}.$$

The first-order condition is

$$-\sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t} - (1-\nu)\sum_{t\geq 0}\mathcal{B}_{g,t}^{2} = 0,$$

and solving for $1-\nu$,

$$1-\nu_g^{\star,R} = -\frac{\sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t}}{\sum_{t\geq 0}\mathcal{B}_{g,t}^{2}} = \mathcal{O}_g,$$

by the definition of the offsetting share (50). Therefore $\nu_g^{\star,R} = 1-\mathcal{O}_g$.

Substituting $1-\nu_g^{\star,R}$ back into to show it attains a minimum,

$$L_g^{R}(\nu_g^{\star,R}) = \tfrac{1}{2}\sum_{t\geq 0}\mathcal{A}_{g,t}^{2} - \tfrac{1}{2}\,\frac{\left(\sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t}\right)^{2}}{\sum_{t\geq 0}\mathcal{B}_{g,t}^{2}} \;\geq\; 0,$$

which is positive since $(\sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t})^2 \leq \sum_{t\geq 0}\mathcal{A}_{g,t}^2 \sum_{t\geq 0}\mathcal{B}_{g,t}^2$.

Next, move to the aggregate problem by multiplying the cohort loss by ω_g and summing across cohorts,

$$\mathcal{L}^R(\nu) = \tfrac{1}{2}\sum_g \omega_g \sum_{t\geq 0}\mathcal{A}_{g,t}^2 + (1-\nu)\sum_g \omega_g \sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t} + \tfrac{1}{2}(1-\nu)^2 \sum_g \omega_g \sum_{t\geq 0}\mathcal{B}_{g,t}^2.$$

The first-order condition gives

$$1-\nu_{\text{agg}}^{\star,R} = -\frac{\sum_g \omega_g \sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t}}{\sum_h \omega_h \sum_{t\geq 0}\mathcal{B}_{h,t}^2}.$$

To write this with the cohort shares, multiply and divide the term of each cohort g in the numerator by $\sum_{t\geq 0}\mathcal{B}_{g,t}^2$,

$$1-\nu_{\text{agg}}^{\star,R} = \sum_g \frac{\omega_g \sum_{t\geq 0}\mathcal{B}_{g,t}^2}{\sum_h \omega_h \sum_{t\geq 0}\mathcal{B}_{h,t}^2}\Big(-\frac{\sum_{t\geq 0}\mathcal{A}_{g,t}\mathcal{B}_{g,t}}{\sum_{t\geq 0}\mathcal{B}_{g,t}^2}\Big) = \sum_g \mathcal{S}_g^R\, \mathcal{O}_g,$$

which imply

$$\nu_{\text{agg}}^{\star,R} = 1 - \sum_g \mathcal{S}_g^R\, \mathcal{O}_g = \sum_g \mathcal{S}_g^R\, \nu_g^{\star,R}$$

as we wanted to prove.

Part 2 (HANK). Fix the realised paths $d\boldsymbol{\pi}$ and $d\boldsymbol{\chi}_g$ of the cost-push shock. By Proposition 4, $\mathcal{W}_{g,t}^H(\nu) = \mathcal{A}_{g,t}^H + (1-\nu)\mathcal{B}_{g,t}^H$, with $\mathcal{A}_{g,t}^H$ and $\mathcal{B}_{g,t}^H$ as in (47). This is the same form as for the HANK case, with each error scaled by the cohort's own Jacobian, so we can repeat the same steps to get the results. The only caveat is that the Jacobians enter the optima, so the HANK optimum differs from the RANK one even on the same inflation and credibility paths. □

Appendix B. Equilibrium and Behavioural Jacobians

B.1. Equilibrium Definition in HANK

DEFINITION A1 (Competitive equilibrium). *Given the initial cohort distributions* $\{D_{g,0}\}_g$ *and exogenous paths* $\{\varepsilon_t^{MP}, \varepsilon_t^{CP}\}$*, an equilibrium is a collection of value and policy functions* $\{V_{g,t}, c_{g,t}, a_{g,t}\}_g$*, distributions* $\{D_{g,t}\}_g$*, prices* $\{w_t, r_t, i_t, \pi_t, \pi_t^w\}$*, cohort credibility sequences* $\{\chi_{g,t}\}_g$*, policies* $\{\tau_t\}$*, and aggregates* $\{C_t, N_t, A_t, Y_t\}$ *such that: (i)* $\{c_{g,t}, a_{g,t}\}$ *solve* (28) *under the diagnostic expectations of Proposition 2 and the*

perceived rate (27)*; (ii)* $\{D_{g,t}\}$ *satisfy* (30)*; (iii) the firm and union optimality conditions* $w_t = \Theta_t$ *and* (32) *hold, with* $Y_t = \Theta_t N_t$ *and* $N_{g,t} = \Gamma_g(N_t)$*; (iv) the government budget and the Taylor rule* (33) *hold; (v) cohort credibility* $\chi_{g,t}$ *evolves according to the law of motion* (A11)*; and (vi) the bond and goods markets clear,*

$$A_t = \sum_g \omega_g A_{g,t} = B, \qquad Y_t = \sum_g \omega_g C_{g,t} + G. \tag{A25}$$

B.2. Sequence-Space Solution Method

This paper uses the sequence space approach introduced by Auclert et al. (2021) due to its effectiveness for solving heterogeneous-agent models when nonlinearities do not arise. In order to obtain the first-order transition dynamics, the economy is written as a set of blocks, each being fed by inputs to produce time paths of aggregate outputs. Define the first-order behaviour of a generic block with output O and input X as its Jacobian matrix

$$\mathcal{J}^{O,X} = \left[\partial O_t / \partial X_s\right]_{t,s\geq 0}, \tag{A26}$$

whose (t, s) entry captures how the date-t output responds to a unit innovation in the date-s input, evaluated at the steady state. Intuitively, column s will be interpreted as the impulse response of O to news, received today, about the date-s value of X. Following Auclert et al. (2021), the household blocks' Jacobians are computed with the fake-news algorithm.

The economy of Section 4.2 contains G cohort household blocks plus the production, fiscal, bond-pricing, Phillips-curve, monetary-policy, diagnostic-expectations, credibility, and market-clearing blocks. Any block other than the g household problems is a simple analytical block that resembles Galí (2015). Each cohort g block is a standard incomplete-markets recursion, and these are the only non-trivial blocks.[15]

[15] I adopt the same solution framework as in Auclert et al. (2021). Given the paths of inputs for every heterogeneous g block, I solve the problem (28) backwards for the marginal value of wealth and the consumption and savings policies $c_{g,t}(a, e)$, $a_{g,t}(a, e)$ using the endogenous-gridpoint method, over a state of idiosyncratic productivity e (evolving on the Markov chain $P(e, e')$) and bond holdings a on a grid. Then, I iterate the cohort distribution $D_{g,t}(\mathrm{d}a, e)$ forward and aggregate the policies. Cohorts follow the same recursion and the only differences are the discount factor β_g, the labour income $y_{g,t}$, and the associated perceived rate $r^{\text{per}}_{g,t}$. Finally, time paths are truncated at a horizon T, and the household Jacobians are computed once at the steady state.

B.3. The Distorted-Beliefs Jacobian

This section makes explicit how to compute the household Jacobians under distorted beliefs. The main idea is to inject the belief distortion into the Euler equation and let the sequence-space solver differentiate the otherwise standard block. The implementation closely mirrors the behavioural-Jacobian representation of Guerreiro (2023) and Shabalina and Tzaawa-Krenzler (2024), in which a household block under non-full-information expectations can be written as the FIRE Jacobian operating on the believed path, plus a correction for forecast revisions.

For the paths of inflation and credibility, a convenient property of the diagnostic beliefs of Proposition 2 is $d\mathbb{E}^{\theta}_{g}[\boldsymbol{\pi}] = \Lambda^{\pi}_{g,t} d\boldsymbol{\pi} + \Lambda^{\chi}_{g,t} d\boldsymbol{\chi}_g$, with $\Lambda^{\pi}_{g,t} = \rho_g[(1+\theta)\mathbb{I} - \theta\rho_g L]$ and $\Lambda^{\chi}_{g,t} = (1-\rho_g)\kappa_g[\mathbb{I} - \theta\rho_g L]$, the operators of (41) stacked over horizons, with L the lag operator. Then, for $X \in \{\pi, \chi\}$, the behavioural Jacobian becomes:

$$\tilde{\mathcal{J}}^{C,X}_{g} = \mathcal{J}^{C,X}_{g} \Lambda^{X}_{g,0} + \sum_{t\geq 1} \mathcal{R}^{C,X}_{g,t} \left(\Lambda^{X}_{g,t} - \Lambda^{X}_{g,t-1}\right), \tag{A27}$$

where $\mathcal{R}^{C,X}_{g,t}$ is the consumption response to a forecast revision at t. Intuitively, the first term is the FIRE Jacobian, while the second term contains all the forecast revisions that distorted beliefs trigger.

In this model the second term is zero. By Proposition 2, the linearised expectation at any horizon $h \geq 0$ is

$$d\mathbb{E}^{\theta}_{g,t+h}[\pi_{t+h+1}] = \rho_g\left[(1+\theta)d\pi_{t+h} - \theta\rho_g\, d\pi_{t+h-1}\right] + (1-\rho_g)\kappa_g\left[d\chi_{g,t+h} - \theta\rho_g\, d\chi_{g,t+h-1}\right],$$

with ρ_g fixed by Assumption 2, θ common and $\kappa_g = -(1-\nu)\Delta_g$ fixed for a given intensity. Hence the entries of $\Lambda^{\pi}_{g,t}$ and $\Lambda^{\chi}_{g,t}$ do not depend on the date at which the expectation is formed, every revision term $\Lambda^{X}_{g,t} - \Lambda^{X}_{g,t-1}$ vanishes, and the behavioural Jacobians collapse to

$$\tilde{\mathcal{J}}^{C,\pi}_{g} = \mathcal{J}^{C,\pi}_{g} \Lambda^{\pi}_{g}, \qquad \tilde{\mathcal{J}}^{C,\chi}_{g} = \mathcal{J}^{C,\pi}_{g} \Lambda^{\chi}_{g}.$$

The second equality uses that the two paths enter the household problem only through the perceived rate (20), $d\mathbf{r}^{\text{per}}_{g} = d\mathbf{i} - d\mathbb{E}^{\theta}_{g}[\boldsymbol{\pi}]$, so the FIRE Jacobian with respect to the expected inflation path is the same whether the expectation moves because of $\boldsymbol{\pi}$ or because of $\boldsymbol{\chi}_g$, namely $\mathcal{J}^{C,\pi}_{g} = \mathcal{J}^{C,\chi}_{g} = -\mathcal{J}^{C,r^{\text{per}}}_{g} = -\mathcal{J}^{C,i}_{g}$. All Jacobians are evaluated at the belief-free steady state, where the perceived and the realised

rate coincide. This is the property that the proof of Proposition 4 uses, and it is why the FIRE Jacobians of the sequence-space solver can be reused under distorted beliefs.

Appendix C. Calibration of the Belief Parameters

This section presents how the parameters of the belief block reported in Table 1 are obtained. The cohort moments ρ_g and $\bar{\pi}_g$ are measured from each cohort's lifetime inflation. The learning parameters ρ_U and b are imported from Hommes and Lustenhouwer (2019), and the diagnostic parameter θ from Carroll (2003). Finally, the response of credibility to inflation in the BoE Inflation Attitudes Survey is used as an untargeted moment to assess the calibration.

C.1. Cohort Moments, ρ_g and $\bar{\pi}_g$

Under Assumption 2, cohort g perceives inflation as a first-order autoregression with persistence ρ_g, and under Assumption 3 its experiential anchor is built on its lived inflation $\bar{\pi}_g$. Both moments are computed from the cohort's lifetime database of quarterly-annualised CPI inflation, namely the $a_g^Q = 4\,(a_g - 16)$ quarters before 2019Q4, with $a_g \in \{20, 30, 40, 50, 60, 70\}$ the cohort age midpoint. The lived inflation $\bar{\pi}_g$ is the mean of that database, and ρ_g is the OLS slope of

$$\pi_s = c_g + \rho_g\, \pi_{s-1} + e_{g,s}, \tag{A28}$$

estimated on the $a_g^Q - 1$ consecutive pairs of that window, which ends in 2019Q3. Panel B of Table 1 reports both moments. They are measured once at 2019Q4, and kept fixed over the horizon of the model.

Two limitations are worth noting. First, the window of the youngest cohort has only 16 quarters, 2015Q4–2019Q3, so its persistence is imprecise. Second, the estimates of the three youngest cohorts, between 0.04 and 0.11 with standard errors between 0.10 and 0.26, are not statistically different from zero. Additionally, the 45–54 cohort, 0.18 (0.09), is significant only at the 5% level; and the two oldest cohorts, 0.56 (0.06) and 0.55 (0.06), are pretty close. Therefore, the jump in ρ_g between the 45–54 and the 55–64 cohorts reflects the entry of the 1970s–80s inflation into the database, not a change in the estimation.

C.2. Learning Parameters, ρ_U and b

Equation (8) ties b to the perceived innovation variance σ_g^2 and to ρ_U. Since σ_g^2 is a perceived variance and not an estimated one, b is not computed from it. Instead, both learning parameters are imported from Hommes and Lustenhouwer (2019), whose measure of credibility has the same logit form as Equation (8) and whose performance measure discounts past squared forecast errors at a geometric rate, as in Equation (9):

- $\rho_U = 0.5$ is their memory parameter.
- $b = 0.25$ is their intensity of choice. They set it to 40,000 for quarterly inflation expressed in decimals. The forecast errors here are in percentage points of annualised inflation, which are 400 times larger, so the same intensity is $b = 40{,}000/400^2 = 0.25$.

C.3. Diagnostic Parameter, θ

The diagnostic parameter is common across cohorts and imported from Carroll (2003). By Equation (13), the diagnostic expectation can be written as

$$\mathbb{E}^{\theta}_{g,t}[\pi_{t+1}] = \mathbb{E}_{g,t-1}[\pi_{t+1}] + (1+\theta)\big(\mathbb{E}_{g,t}[\pi_{t+1}] - \mathbb{E}_{g,t-1}[\pi_{t+1}]\big),$$

so it absorbs a fraction $1+\theta$ of the forecast revision on impact. Carroll (2003) estimates that a fraction $\lambda = 0.27$ of households update their inflation expectation each quarter, so the mean household expectation absorbs the fraction λ of the news on impact. Equating the two gives

$$\theta = \lambda - 1 = -0.73.$$

The mapping is exact on impact only. After the first quarter the diagnostic expectation absorbs the rest of the revision, while under sticky information the revision keeps being absorbed gradually. The estimate of Carroll (2003) comes from the Michigan Survey of US households. It is used here because the BoE survey is a repeated cross-section, which rules out estimating an updating probability at the individual level.

C.4. Lived Inflation and Credibility

Since the learning parameters are imported, the response of credibility to inflation in the data is not used in the calibration and serves as a check on it. The measure of credibility is the aggregate anchoring share χ_t, the share of respondents who answer question Q18 of the BoE Inflation Attitudes Survey "very" or "fairly" confident that inflation will be close to the target (Section D). The question is asked in the February wave of every year since 2011 and, in addition, in 2025Q3, so the sample has 17 waves, 2011Q1–2026Q1. I regress aggregate share χ_t on its value in the previous wave and on the inflation gap of the quarter of the previous wave,

$$\chi_t = c + \rho_y \chi_{t-1} - \eta_y (\pi_{t-1} - \pi^*) + \varepsilon_t, \tag{A29}$$

where $t-1$ denotes the previous wave, π_t is the quarterly-annualised CPI inflation of the quarter of wave t in percentage points and $\pi^* = 2$, which gives 16 observations. The equation is estimated by OLS. The estimated parameters are $\rho_y = 0.317$ (0.220) and $\eta_y = 0.0163$ (0.0070), with $R^2 = 0.45$.

The implied long-run response of the anchoring share to a permanent one-percentage-point inflation gap is $d\chi/d\pi|_{LR} = -\eta_y/(1-\rho_y) = -0.024$.

The model counterpart follows from Equation (A11). A permanent unit inflation gap is a permanent quasi-innovation of $1-\rho_g$, so at $\nu = 0$

$$\left.\frac{d\chi}{d\pi}\right|_{LR} = -\frac{1}{100}\sum_g \omega_g \frac{(1-\rho_g)\,\xi_g}{1-\rho_U} = -0.039, \tag{A30}$$

The imported values of ρ_U and b in Therefore, Table 1 shows a de-anchoring larger than in data, 1.3–1.6 times, but within the 95% interval of the static estimate, $[-0.045,\ -0.015]$. Given 16 observations, most of them from the 2022–2024 episode, the check is indicative rather than a test.

Appendix D. Empirical Data

This appendix shows in detail how the empirical variables mentioned in Section 2, and used throughout the paper, are constructed. It covers the survey questions and samples, the inflation history, the wealth and income data behind the HANK calibration, and some descriptive statistics.

D.1. Survey Questions

This section reproduces the questions of the Bank of England / Ipsos Inflation Attitudes Survey (IAS) that are used along the empirical analysis of Section 2. The survey is administered to a repeated cross-section of UK adults. Each response has assigned a post-stratification weight, used as the observation weight throughout. In what follows, the relevant sections of the survey for this paper are presented.

D.1.1. Coverage

The IAS provides quarterly data, collected every February, May, August and November since 2001Q1 and it is currently conducted by Ipsos. Each wave is a new, nationally representative sample of UK adults, so the data form a repeated cross-section rather than a panel. Cohorts are therefore not followed at the individual level. Until November 2021 the survey was collected in-home, face-to-face, on a random-location quota sample of UK adults aged 16 and over. Since February 2022 Ipsos runs it on an online format of about 2,000–2,200 UK adults aged 16–75, with quotas on age, gender, region and working status. The 2022 redesign therefore changes the mode, the upper age limit, and the coding of the age and income bands (Section D.1).

Four windows are used in the paper, one per empirical object.

a. *Forecast errors and revisions* (Fact 1, Questions 2a and 2b). Quarterly waves 2010Q1–2024Q4, 60 waves per cohort, with the three cohorts 15–44, 45–64 and 65+, labelled ≤ 45, 45–65 and 65+ in the figures following Gennaioli et al. (2024). Every figure of the paper uses these three cohorts, obtained by averaging the six age bands within bin and then aggregating the bins with their survey weights.
b. *Credibility* (Fact 2, Question 18). The question is asked in the February wave of every year since 2011 and, in addition, in the August 2025 wave. The figures use the 16 waves 2011Q1–2025Q3 and the regression of Appendix C.4 the 17 waves 2011Q1–2026Q1.
c. *Forecast cues* (Fact 3, Question 2aiv). Available in only six waves, 2017Q1, 2023Q1, 2024Q1, 2025Q1, 2025Q3 and 2026Q1, so this is a separate cross-section from the other windows. The income subgroups use the post-2022 income bands.
d. *Calibration moments*. The cohort inflation experiences behind ρ_g and $\bar{\pi}_g$ are the a_g^Q quarters before 2019Q4, i.e. the window ending in 2019Q3 (Appendix C.1); the cohort income profile uses the waves up to 2021Q4, before the 2022 redesign of the income bands, and the wealth data are the ONS round of April 2018–March

2020 (Section D.3).

Throughout the paper, the post-stratification weights provided by the survey are used. Thus, each cross-section aligns with the UK population profile by age, gender, working status and region. Next, I reproduce the particular questions of interest used throughout.

D.1.2. Demographics

Age [age].. Six ordinal bands that index the cohort $g \in \{1, \ldots, 6\}$:

- 1 = 15–24,
- 2 = 25–34,
- 3 = 35–44,
- 4 = 45–54,
- 5 = 55–64,
- 6 = 65+.

The age band fixes the cohort lifetime inflation database, with memory beginning at age 16 following Gennaioli et al. (2024) (Section D.2). From 2022Q1 the online samples ages 16–75 and stores age under a different coding; its codes 2, 3, 4, 5, 7 and 8 are recoded onto the original bands 1 to 6, so the first band covers ages 16–24 and the last one ages 65–75 in the post-2022 waves.

Gross Household Income [income].. Annual gross household income band. The coding changes in 2016 and again in 2022 as follows:

- 1 = <£9,500 *[removed 2022 Feb]*,
- 2 = £9,500–17,499 *[removed 2022 Feb]*,
- 3 = £17,500–24,999 *[removed 2022 Feb]*,
- 4 = >£25,000 *[removed 2016 Feb]*,
- 5 = £25,000–39,999 *[added 2016 Feb, removed 2022 Feb]*,
- 6 = >£40,000 *[added 2016 Feb, removed 2022 Feb]*,
- 7–11 = post-2022 bands (<£9,999; £10,000–19,999; £20,000–34,999; £35,000–44,999; >£45,000) *[added 2022 Feb]*,
- 12 = prefer not to answer *[added 2022 Feb]*.

The pre-2022 bands are collapsed into three groups when income enters as a control: LOW = $\{1, 2\}$ (<£17,500), MID = $\{3\}$ (£17,500–24,999), HIGH = $\{4, 5, 6\}$ (>£25,000). The cohort income profile of the HANK model uses the pre-2022 coding (Section D.3). The income subgroups of Fact 3 use the post-2022 bands, since Question 2aiv exists only in 2017Q1 and from 2023Q1 onwards: LOW = $\{7, 8\}$ (<£20,000),

MID = $\{9\}$ (£20,000–34,999) and HIGH = $\{10, 11\}$ (>£35,000), with "prefer not to answer" excluded and the 2017Q1 wave left out of the income rows because its bands are not comparable.

Education [`educ`], housing tenure [`tenure`] and Question 11.. These variables define the subgroups of Fact 3 and the controls of the credibility regressions. Education has three levels, GCSE, A-level and degree. Housing characteristic is recorded as owned outright, mortgage, council rent and other, and the last two are gathered as "rent". Question 11 asks "Each month a group of people meets to set Britain's basic interest rate level. Do you know what this group is?", with the options Monetary Policy Committee, Bank of England, the Government, the Treasury, Parliament, other, and don't know. Respondents who answer Monetary Policy Committee or Bank of England are classified as knowing who sets Bank Rate, those who name another body as not knowing, and "don't know" is excluded from this.

D.1.3. Perceived Inflation

Question 1 — perceived inflation over the past 12 months [`q1`] .

- *"Which of these options best describes how prices have changed over the last 12 months?"*

This question captures the household's perception of realised inflation over the preceding year, with the following options to pick from:

- 1 = gone down,
- 2 = not changed,
- 3 = up by 1% or less,
- 4 = up by 1–2%,
- 5 = up by 2–3%,
- 6 = up by 3–4%,
- 7 = up by 4–5%,
- 8 = up by 5% or more,
- 9 = don't know.

I map each band to its midpoint in percent for the codes 3–7. The answer "not changed" (code 2) is set to zero and "gone down" (code 1) to −0.5%, while the top band "up by 5% or more" (code 8) is top-coded at 6.5%. The "don't know" code (9)

is set to a missing value. This yields the perceived twelve-month inflation $\pi_{i,t}^{\text{perc}}$, used in Fact A2 of Appendix E, with coverage 2001Q1 – 2026Q1.

D.1.4. Inflation Expectations

Question 2a — one-year ahead inflation expectation [`q2a_agg1`] .

- *"How much would you expect prices in the shops generally to change over the next 12 months?".*

Question `q2` offers a scale (go down / not change / up by 1% bands up to "5% or more"), with follow-ups capturing larger rises and falls and a further >10% extension. The BoE harmonises these into the aggregate variable `q2a_agg1`, with the following options to pick from:

- 1 = Down by 1% or less,
- 2 = Down by 1–2%,
- 3 = Down by 2–3%,
- 4 = Down by 3–4%,
- 5 = Down by 4–5%,
- 6 = Down by 5% or more,
- 7 = not changed,
- 8 = up by 1% or less,
- 9 = up by 1–2%,
- 10 = up by 2–3%,
- 11 = up by 3–4%,
- 12 = up by 4–5%,
- 13 = up by 5–6%,
- 14 = up by 6–7%,
- 15 = up by 7–8%,
- 16 = up by 8–9%,
- 17 = up by 9–10%,
- 18 = up by 10% or more,
- 19 = no idea,
- 20 = up by 10–11%,
- 21 = up by 11–12%,
- 22 = up by 12–13%,
- 23 = up by 13–14%,
- 24 = up by 14–15%,
- 25 = up by 15% or more *(codes 20–25 added 2022 May)*.

To process this question, I map each closed band to its midpoint in percent per year, assign the open bands the value half a point beyond their threshold (−5.5, 10.5 and 15.5 for codes 6, 18 and 25), and set the "no idea" code (19) to a missing value. The weighted mean within the cohort and wave is the one-year expectation $\mathbb{E}_{g,t}[\pi_{t,t+3}]$ of Fact 1, written $\mathbb{E}_g[\pi_{t+4}]$ in Table A1.

Question 2b — inflation expectation for the year after next [`q2b_agg1`] .

- *"And how about prices the year after next? How much would you expect prices in the shops generally to change between [month, year +2] and [month, year +3]?"*

Recorded and processed as `q2a_agg1`. Together with Question 2a it builds the fixed-event revisions and forecast errors of Fact 1: the expectation a cohort held one year earlier for the twelve months that Question 2a covers today comes from Question 2b of the wave four quarters before. The realised inflation that closes the forecast error is the mean of the four quarterly annualised CPI rates over quarters t to $t + 3$, so the last usable wave is 2024Q4.

Question 2c — five-years ahead inflation expectation [`q2c_agg1`] .

- *"And how about the longer term, say in five years time? How much would you expect prices in the shops generally to change over a year then?"*

Recorded and following the same structure as in `q2a_agg1` above. Similarly, to process it a midpoint map is used and "no idea" is set to missing value. This variable is labeled as $\mathbb{E}_g[\pi_{t+20}]$.

Question 2aiii — recent revision of inflation expectations [`q2aiii`]. .

- *"How has your expectation [of expected price change over the next 12 months] changed over the past 12 months?"*

Asked only to those who gave a numeric answer to Question 2a. The options are:

- 1 = Risen a lot,
- 2 = Risen a little,
- 3 = Unchanged,

- 4 = Fallen a little,
- 5 = Fallen a lot,
- 6 = No idea.

It is the filter of Question 2aiv and defines the "direction of the revision" subgroups of Fact 3 (up if 1 or 2 was chosen, down if 4 or 5).

Question 2aiv — self-reported forecast cues `[q2aiv1–q2aiv13]`.

- *"What were the most important factors that led you [to revise / not revise] your expectation of prices in the shops over the next 12 months? Please select up to 4."*

First asked 2017 Feb to those who answered 1–5 in Question 2aiii. The response is stored as thirteen binary indicators $d_{i,t}^{(k)} \in \{0, 1\}$:

a. How prices have changed in the shops recently, over the last 12 months;
b. How prices have changed in the shops, on average, over the longer term, i.e. the last few years;
c. Reports of current inflation in the media;
d. Discussion of the prospects for inflation in the media;
e. The level of interest rates;
f. The inflation target set by the government;
g. The current strength of the UK economy;
h. Expectations about how economic conditions in the UK are likely to evolve;
i. The level of the exchange rate (the value of sterling);
j. Other factors;
k. *(unused code)*;
l. How the price of energy (e.g. gas, electricity or liquid fuel bills) has changed recently, over the last 12 months *[added 2023 Feb]*;
m. How the price of petrol or diesel has changed recently, over the last 12 months *[added 2023 Feb]*.

The sample for 2aiv has only six waves: 2017Q1, 2023Q1, 2024Q1, 2025Q1, 2025Q3 and 2026Q1, so Fact 3 is a separate cross-section from the other windows. The paper uses the two cues the central bank controls through its communication: the inflation target, option 6, whose weighted share is $\hat{\mu}_T$, and the level of interest rates, option 5, whose weighted share is $\hat{\mu}_I$. The difference $\hat{\mu}_I - \hat{\mu}_T$ is computed within respondent, so both shares come from the same sample. Its standard errors are clustered by wave × age-band cell.

D.1.5. Behavioural response and credibility

Question 17 — actions expected to be taken [q17_1–q17_8].

- *"Which, if any, of the following actions are you taking, or planning to take, in the light of your expectations of price changes over the next twelve months? You can choose up to three."*

This question is stored as eight binary indicators:

a. Bring forward major purchases such as furniture or electrical goods;
b. Cut back spending and save more;
c. Shop around more for better value goods and services;
d. Push for increased pay with current employer;
e. Look to increase income in other ways (e.g. change jobs, take on a second job, work more hours with current employer);
f. Move savings out of banks or building societies into other assets such as shares, bonds, housing or gold;
g. Take no action;
h. None / Don't know.

Fact A1 of Appendix E uses the two options that map to the consumption–saving margin, as weighted shares by cohort and wave:

a. `q17_1` (bring forward purchases — intertemporal substitution),
b. `q17_2` (cut back spending and save more — precautionary).

Question 17 is asked in the same waves as Question 18, once a year in February since 2011 and in 2025Q3. The share is the weighted fraction of all respondents who received the question and selected the option.

Question 18 — anchoring / credibility [q18].

- *"How confident are you that inflation will be close to the 2% target set by the government in two to three years' time? By close I mean that inflation will be between 1% and 3%."*

Available options:

- 1 = Very confident (better than 50% chance),
- 2 = Fairly confident (around 50% chance),
- 3 = Not very confident (around 25% chance),
- 4 = Not at all confident (less than 25% chance),
- 5 = Don't know.

Note. The first wave with this question, 2011Q1, used an earlier "chance" wording without a "don't know" option, mapped onto the same four ordinal levels: 6 = very good chance (>75%); 7 = good chance (50–75%); 8 = some chance (25–50%); 9 = not much chance (<25%). The current wording and the "don't know" option were introduced in 2012Q1. To process this question, I drop "don't know" (code 5) from the set of valid answers $\mathcal{V}_{g,t}$ and map the remaining codes to the four ordinal levels (with the old-wave codes 6–9 mapped to 1–4). The credibility measure of the paper, Equation (2), is the weighted share of the cohort that answers "not very" or "not at all" confident, $1 - \hat{\chi}_{g,t}$, so that $\hat{\chi}_{g,t}$ is the share answering "very" or "fairly" confident. The shares are computed within each wave and age band with the survey weights. The cohorts of the figures are the averaged bands and fewer than 30 valid answers are dropped.

Question 9b — perceived transmission of interest rates [`q9b`].

- *"A rise in interest rates would make prices in the high street rise more slowly in the medium term (say a year or two)."*

Respondents state whether they agree with the sentence on a five-point scale from "strongly agree" to "strongly disagree", plus "don't know". Fact A2 plots, by cohort and wave, the weighted share who agree or strongly agree and the share who disagree or strongly disagree, both excluding "don't know" from the denominator, over 2001Q1–2026Q1.

D.2. Lived Inflation

Lived inflation is calculated using the official quarterly UK inflation rate provided by the Office for National Statistics. The constructed series spans from 1949Q2 to 2026Q1 and covers every cohort's lifetime memory window. First, I combine the monthly historical modelled index `JFDN` (*CPI HISTORICAL MODELLED INDEX 00, 1965=100*, Jan 1949 – Dec 1988) with the recent index `D7BT` (*CPI INDEX 00, 2015=100*, monthly 1988 – Mar 2026). The two indices are linked at January 1988. The historical index is rescaled by the ratio of the two series in that month and used up to December 1987, and the current index is used from January 1988 onwards.

In line with Gennaioli et al. (2024), the series is the annualised quarter-on-quarter growth rate of the end-of-quarter CPI level, $\pi_t = ((\mathrm{CPI}_t/\mathrm{CPI}_{t-1})^4 - 1) \times 100$.

For every individual i in cohort g surveyed in quarter w, the per-cohort historical database is the sequence of quarterly annualised inflation observations $\Pi^{\text{hist}}_{g,w} = \{\pi_{w-2-s}\}_{s=0}^{a_g^Q-1}$, with $a_g^Q = 4\,(a_g - 16)$ quarters of memory (memory starting at age 16 following Gennaioli et al. (2024)), and $a_g \in \{20, 30, 40, 50, 60, 70\}$ the cohort age midpoint, so the database ends one quarter before the most recent completed quarter at the time of the survey. The lived inflation $\bar{\pi}_g$ plotted in Facts 2 and A2 is the mean of this database at each wave, averaged over the age bands of the cohort with their survey weights. The calibration moments ρ_g and $\bar{\pi}_g$ of the model use the database of the a_g^Q quarters ending in 2019Q3, the last complete quarter before 2019Q4, and are kept fixed (Appendix C.1).

D.3. Wealth, Income and Discount Factors

The lifecycle HANK model of Section 4.2 needs, for each cohort $g \in \{1, \ldots, 6\}$, a labour income level and a wealth level. Both are measured at the cohort level, matching the survey age bands, and they pin down the cohort discount factors β_g. The marginal propensities to consume are then outcomes of the solved model, not inputs.

Income. The cohort income profile comes from the IAS income question (Section D.1), using the pre-2022 coding over the waves 2001Q1–2021Q4. The closed bands below £25,000 are valued at their midpoints (£4,750, £13,500 and £21,250); the band £25,000–39,999 at £45,000 and the open bands at £32,500 (>£25,000, until 2015) and £65,000 (>£40,000, from 2016). The survey-weighted mean by age band gives the cohort income Y_g. Dividing by the population-weighted mean income gives the relative income profile that scales the idiosyncratic productivity of each cohort. Thus, all cohorts share the same income risk but differ in income level.

Wealth. Per-cohort wealth is the mean total individual wealth by age band of the ONS *Wealth and Assets Survey* (WAS), round April 2018 – March 2020, the latest available and the one that matches the memory database frozen at 2019Q4. The published five-year age bands are aggregated to the six IAS cohorts using the ONS population frequencies as weights, delivering the wealth profile W_g. The mean, rather than the median, is the object that clears the asset market in the steady state. The WAS covers Great Britain and individuals aged 16 and over; total wealth

is the sum of property, financial, physical and private pension wealth. To keep consistency with the age groups, the bands 16–19 and 20–24 form the first cohort and the six bands from 65–69 to 90 and over form the last one.

Discount Factors. Within each cohort, households are of two types that differ only in patience, with discount factors β_g. The wealth-to-income ratios $(W/Y)_g$ calibrate the patient discount factors through a shape-and-level decomposition, $\beta_g = \bar{\beta} + \delta_g$. The shape δ_g is proportional to the log wealth-to-income ratio centred at its population-weighted mean, scaled so that the distance between the highest and the lowest δ_g is 0.003.

D.4. Descriptive Statistics

TABLE A1. **Data Sources and Descriptive Statistics**

Variable	N	Mean	SD	Min	Max
Panel A: Inflation Expectations					
BoE/Ipsos Inflation Attitudes Survey; individual responses, survey-weighted					
quarterly, 2001Q1–2026Q1 (95 waves, n = 250,837)					
$\mathbb{E}_i[\pi_{t+4}]$ (Q2a, next 12 months, %)	155,882	3.65	3.29	−5.50	15.50
$\mathbb{E}_i[\pi_{t+8}]$ (Q2b, year after next, %)	133,047	3.22	3.10	−5.50	15.50
$\mathbb{E}_i[\pi_{t+20}]$ (Q2c, five years ahead, %)	123,622	3.82	3.48	−5.50	15.50
survey weight w_i	250,837	1.00	0.4528	0.0011	13.49
Panel B: Credibility					
BoE IAS Q18; cells wave × cohort (3 cohorts), 16 waves 2011Q1–2025Q3 (48 cells)					
aggregate χ_t over the 17 Q18 waves, 2011Q1–2026Q1 (February waves and 2025Q3)					
$\hat{\chi}_{g,t}$ (share very or fairly confident, cell)	48	0.4212	0.1121	0.1929	0.6191
$1-\hat{\chi}_{g,t}$ (lack of credibility, cell)	48	0.5788	0.1121	0.3809	0.8071
respondents per cell	48	582.4	203.0	220.0	1,036.0
χ_t (aggregate anchoring share)	17	0.4226	0.1020	0.2527	0.5852
$\lvert\Delta\chi_t\rvert$ (abs. wave-to-wave change)	16	0.0689	0.0541	0.0078	0.2142
Panel C: Inflation Experience Series					
ONS CPI, series JFDN spliced to D7BT in January 1988; quarterly, annualised					
1949Q2–2026Q1 (308 quarters)					
π_t (annualised CPI inflation, %)	308	4.73	5.49	−6.95	43.10
Panel D: Cohort Wealth & Income					
ONS Wealth and Assets Survey, mean total wealth per adult, April 2018–March 2020 (Great Britain)					
BoE IAS income bands, waves ≤ 2021Q4; six age cohorts					
W_g (mean individual wealth, £000)	6	276.1	217.4	22.45	548.1
Y_g (mean cohort income, £000)	6	23.92	3.94	17.60	27.71
$(W/Y)_g$ (wealth-to-income ratio)	6	12.21	10.80	1.07	27.39

Appendix E. Additional Empirical Facts

This appendix documents two additional facts important for the transmission mechanism of beliefs. Fact A1 shows that when households expect more inflation they report changing the timing of their consumption, on both the spending and

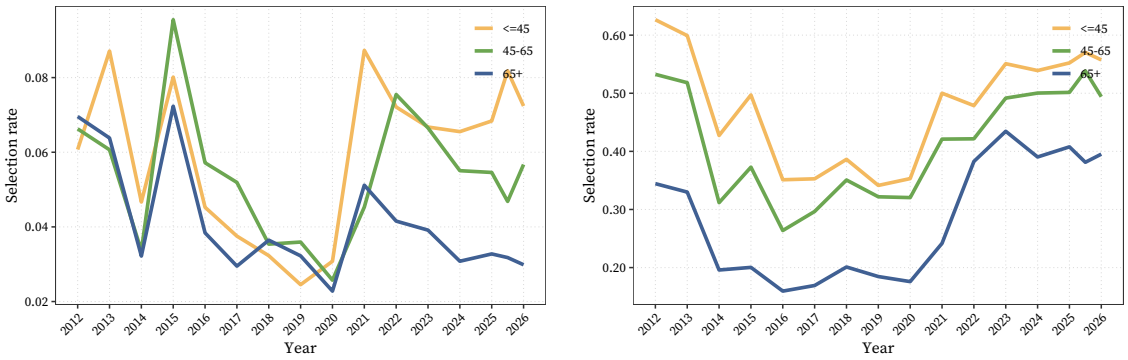
the saving margin. This provides evidence of an intertemporal Fisher effect in place. Fact A2 shows that the same cohorts that lived through higher inflation perceive higher inflation today and, since the 2021–22 surge, have become more sceptical that higher interest rates reduce prices.

E.1. Fact A1. Planned-Consumption Pass-Through.

"When expected inflation rises, a larger share of households report higher intentions to bring forward major purchases and to cut back and save more."

The distorted beliefs channel stated in the paper is relevant only if revisions in expectations trigger changes in consumption decisions. The standard mechanism operates through the cohort Euler equation: a household that expects higher inflation perceives a lower real interest rate, which increases the incentive to bring consumption forward. Therefore, in theory, heterogeneity in expectations would directly drive heterogeneity in consumption timing. The link between inflation expectations and spending has been examined empirically by Coibion et al. (2022), who find that households act on their inflation expectations when making spending and saving decisions.

FIGURE A1. **Consumption Responses to Inflation Expectations**



A. Bring forward major purchases.

B. Cut back spending and save more.

Note: Weighted share of respondents in each cohort selecting the action in BoE IAS question Q17.

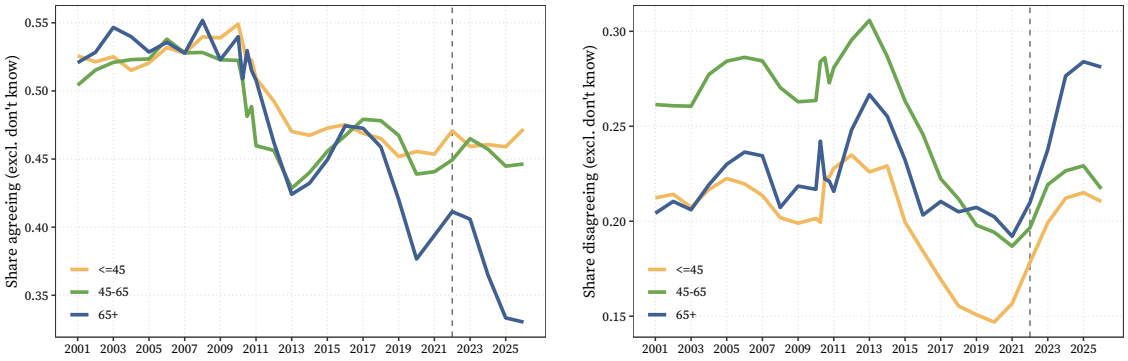
This empirical regularity may also be examined annually in the BoE IAS. Question 17 asks which actions households are planning to take in response to their twelve-month-ahead inflation expectations. Among the potential responses, I focus on the two options related to consumption behavior: bringing forward major

purchases (the intertemporal-substitution response) and cutting back spending to save more. Figure A1 depicts the weighted share of each cohort selecting each intended action. Importantly, both margins move with the inflation cycle: the shares fall through the low-inflation 2010s and rise sharply during the 2021–2023 surge. The responses are also heterogeneous across cohorts: during episodes of higher inflation, younger and middle-aged cohorts tend to report higher intentions to adjust consumption on both margins than older cohorts do.

E.2. Fact A2. The Perceived Channel.

"Older cohorts that have lived through higher inflationary episodes perceive higher inflation. Besides, since the 2021–2022 inflationary shock, they have become less convinced that higher interest rates restrain price increases."

FIGURE A2. **Perceived Monetary Transmission by Cohort**



A. Share agreeing. B. Share disagreeing.

Note: Asked in question Q9b. Both shares exclude "don't know" from the denominator. Included waves 2001–2026.

The final fact provides evidence on the influence of households perception about the transmission of monetary policy and their perception of the inflation rate. To document the change in perceived transmission, the survey asks whether households agree with the following statement: "A rise in interest rates would make prices in the high street rise more slowly in the medium term (say a year or two)." Panel A2A plots the weighted share who answered "agree" or "strongly agree," while Panel A2B plots the share who answered "disagree" or "strongly disagree." Both panels mirror each other. Before 2022, the proportion of households reporting agreement or disagreement moves in parallel across cohorts. However, after the

2021–2022 inflationary shock, the oldest cohort reports more scepticism about whether higher rates would reduce inflation.

Next, turning to perceived inflation, Question 1 of the survey asks about the households' perception of how prices have changed over the last twelve months. Figure A3 plots the inflation experiences against the perceived inflation rate for each cohort and reveals that older cohorts with a higher experienced inflation report higher perceived inflation. This result aligns with Malmendier and Nagel (2016), in the sense that the realized inflation history a cohort has lived through shapes how it perceives current prices. It differs, however, from attention-based models of perceived inflation such as Francis (2026), in which perceptions diverge from official inflation because households overweight salient or volatile goods.

FIGURE A3. **Perceived Inflation and Inflation Memories**

