Distorted Beliefs, Uninsurable Risk and the Transmission Mechanism*

Ester Faia
Goethe University Frankfurt and CEPR

First draft: June 2026; this draft: September 14, 2026

Do no circulate without author's permission

Abstract

The growing availability of survey evidence has shown that belief formation can be a crucial part of households consumption and labor decisions; in turn those may affect the policy transmission mechanism. Using several portions of subjectives belief available in the Survey of Consumer Expectations I estimate a generalized tilt to belief formation, that nests components present in all leading theories of belief formation. I then augment a model with uninsurable risk and incomplete markets with the exact estimated tilt: the households hold an exponentially tilted subjective measure over their own income risk, and the same tilt affects a Bayesian belief update over the dynamic. Coherently and consistently, households form expectations at all horizon with the same operators. The distorted belief helps to reconcile the micro elasticities and two stylized facts about the aggregate response to policy shocks: a generalized hump shaped dynamic and the asymmetric response (recessions hit harder than booms).

Keywords: *distorted beliefs, uninsurable risk, incomplete markets, transmission mechanism.*

JEL Classification: *F41, F13, F44, D57, D52, D24.*

1 Introduction

Heterogeneous agent models have become the standard laboratory for studying how monetary and fiscal policy reach households. The wealth of survey and experimental data has also increased awarness that belief formation process, and the related distortions, affect households consumption-saving and employment decisions. Through this they would in turn affect the transmission mechanism to policy shocks. The shape that this belief formation processes may take can however vary depending on the theories, thereby driving then conclusions about the transmission mechanism. In this paper I use several portions of the Survey of Consumer Expectations to estimate a generalized tilt to belief formation, that nests components present in all leading theories of belief formation. I then augment a model with uninsurable risk and

*Correspondence to: Ester Faia, Chair in Monetary and Fiscal Policy, Goethe University Frankfurt, Theodor W. Adorno Platz 3, 60323, Frankfurt am Main, Germany. E-mail: faia@wiwi.uni-frankfurt.de. Webpage: http://www.wiwi.uni-frankfurt.de/profs/faia/.

incomplete markets with the exact estimated tilt. The belief process has two components, governed by the same operator. In the stationary equilibrium households hold an exponentially tilted subjective measure over their own idiosyncratic income risk: plenty of evidence document households' limited capacity to foresee their earning progression. Based on the estimated process, they overweight states in proportion to how costly those states are, but remain anchored to some reference level. Over the business cycle they do not know which income-risk regime the economy is in and learn about it, with the same tilt applied to the posterior over regimes. The first component shapes precautionary saving, the wealth distribution and the marginal propensity to consume. The second is a belief distortion, which moves in calendar time, more directly affects the transmission of macro shocks.

Our first contribution is the estimate of a generalized belief formation process. Every theory of distorted beliefs is a reweighting of the true probabilities. They differ in whether the weights track how costly a state is, and in whether the reweighting starts from the household's true odds or from an anchor (similar generalized specification is adopted in Bassanin et al. [2021]). Both differences show up in the gap between what households say and what happens to them, and one regression recovers them. We estimate that wedge in the New York Fed Survey of Consumer Expectations, on 69,806 observations covering 11,233 households. The design exploits the twelve-month rotating panel: each household reports a probability of losing its job and is then observed long enough to see whether it did, so the subjective and the objective measure come from the same population and the same instrument. From the estimates households are pessimistic about their own employment risk in every liquidity quintile, with t statistics between 4.5 and 11.3, and the sign is confirmed independently on the job-finding margin. A single specification separates the two components. Cost-ordered tilting moves subjective probabilities away from the truth by a constant amount, whatever the odds; anchoring instead ties them partly to a reference that does not vary with the household's circumstances, so they respond less than fully to the household's actual risk. We find both, with the anchor larger. A further test, asking whether stated beliefs predict spending once true risk is controlled for, shows that most of the anchoring is survey response behaviour rather than a genuine feature of beliefs, which is what keeps it from dominating the calibration.

We then embed the estimated process in a heterogeneous agent New Keynesian model with incomplete markets and borrowing constraints, and we use the survey estimates to set the belief block rather than choosing it to make the model fit. We then examine how the model departs from the comparable benchamrk model with uninsurable risk, focusing mostly in how to reocncile micro and macro elasticities. The impact consumption response to a twenty-five basis point easing is five times smaller, falling from 0.18 to 0.03 percent. The peak moves from

impact to the following quarter. And transmission is asymmetric: a contraction of identical size produces a cumulative response twenty percent larger than an easing, because the tilt on the posterior flips sign with the shock, so a pessimist underweights good news and overweights bad news. No symmetric friction delivers the third of these. On the micro side the marginal propensity to consume rises from 0.096 to 0.165, in the direction the recent evidence points, and nothing was targeted at it.

Two robustness exercises allow to sharpen the channels. First, we show that adding capital with investment adjustment costs and variable utilisation, the canonical device for generating hump-shaped responses, does not change the transmission. Second, introducing the same belief process in a representative agent version of the model separates the two frictions cleanly. The hump and the dampening survive without heterogeneity, because learning is a transformation of the impulse response and is active in any model with a consumption function. The asymmetry does not: the ratio of contraction to expansion falls from 1.20 to 1.00. The reason is quantitative. Asymmetry requires the tilt on the regime posterior to be non-negligible, and that tilt is ϑ times the value of the regime difference. A representative agent has only the aggregate consumption path to put there, worth 3.7×10^{-3} utils, against 1.1 utils for a recession-sized rise in idiosyncratic income risk. In this case then uninsurable risk is the element that chiefly amplifies the distortion from incomplete markets.

We then extends the exploration of the implications of this belief distortion for the transmission mechanism to the fiscal side. The same friction that dampens monetary policy amplifies fiscal policy. Under a balanced budget the impact multiplier of a government spending shock rises from 0.75 to 0.94 once beliefs are introduced, and under a rule in which spending is debt-financed and taxes adjust gradually, from 0.70 to 0.85. The sign reversal is not a coincidence of parameters but follows from how the two policies work. Monetary policy operates by inducing a private response, so a friction that delays that response weakens it. Fiscal policy operates despite an offsetting private response, so the same delay postpones the crowding-out and strengthens it.

The financing rule matters in a way that is again specific to heterogeneous agents. Deficit financing raises the cumulative multiplier, as postponing taxes should, but lowers the impact multiplier. The deficit increases the supply of government debt, and with the real rate pinned by the policy rule the asset market can only clear if households absorb it by saving more. Consumption therefore falls more on impact under the rule even though taxes have not yet risen, a channel with no counterpart in a representative agent model.

Taken together the exercises provide a reasonably sharp conclusion. Uninsurable risk alone delivers the known amplification of aggregate demand. Belief distortions alone deliver dampening and a delayed peak, and would do so in a representative agent model. What requires both is

the asymmetry of monetary transmission, and the reason is that the belief distortion is ordered by the value of states, so it is quantitatively irrelevant unless the states differ by an amount that only uninsurable risk can generate. The two frictions are complements in the transmission mechanism rather than substitutes, and the direction of their joint effect depends on the policy: they weaken monetary policy and strengthen fiscal policy, for reasons that are the same reason seen from two sides.

The paper proceeds as follows. Section 2 estimates the belief process. section 3 states the model and Section 3.2 the calibration. Section 4 reports monetary transmission and the capital extension, and Section 5 the representative agent comparison and the fiscal experiments. Section 6 concludes.

Related Literature: To be Completed. There is growing literature that departs from rational expectations inside a heterogeneous agent model. Farhi and Werning [2019] introduce level-k thinking into an incomplete markets economy and show that bounded rationality and market incompleteness are complements: each attenuates forward guidance, and together they attenuate it far more than either alone. Our finding on the asymmetry has the same complementarity flavour, though it runs through a different object. In their setting the interaction works through the horizon over which households reason; here it works through the size of the value differences the distortion is ordered by, which is why the effect disappears when uninsurable risk is removed. Level-k reasoning is a truncation of the expectation operator rather than a reweighting of states, so it is not nested by (1) and is therefore not one of the alternatives our estimates can rank. Early on, Carroll [2003] proposed that households absorb macroeconomic news slowly, as though it spread through a population by contagion, and documented the resulting sluggishness in survey expectations. Carroll et al. [2020] embed the same idea in a buffer-stock model and show that sticky expectations reconcile the excess smoothness of aggregate consumption with high micro marginal propensities to consume. Auclert et al. [2020] estimate a heterogeneous agent New Keynesian model with sticky household expectations, matching both micro evidence on marginal propensities to consume and the macro impulse response to a monetary shock, and find that inattention is essential for reconciling micro jumps with macro humps. Our learning block plays a related role in generating a delayed peak,. Two key differences are worth stating. Their inattention parameter is estimated from aggregate impulse responses, while ours is estimated from the gap between subjective and realised transitions in household data and then imposed. And sticky information is inert in a stationary environment, so it cannot produce the steady-state pessimism we measure, which is why our specification carries a value-ordered component alongside it. Our learning process is also value-weighted, while their infrequent update is governed by an exogenous poisson probability. At last, Pfäuti and Seyrich [2026] build

a behavioural heterogeneous agent model in which households cognitively discount aggregate news, and show that the amplification of monetary and fiscal policy disappears when either heterogeneity or bounded rationality is removed. Our asymmetry result has the same complementary structure. The difference is that their friction applies to beliefs about aggregates and is calibrated, while ours applies to beliefs about own income risk and is estimated, which is also why their interaction amplifies contemporaneous monetary policy where ours dampens it.

The belief specification adopted here has a direct antecedent in Bassanin et al. [2021], where state-contingent ambiguity attitudes under kinked multiplier preferences generate left-skewed beliefs in recessions and right-skewed beliefs in booms, and the attitude process is estimated by GMM with latent value functions. The asymmetry of policy transmission found below is the macroeconomic counterpart of the asymmetry in beliefs documented there. Two things differ: the distortion applies to households' own income risk rather than to collateral values, and it is estimated from subjective probabilities in household survey data rather than from asset price and leverage moments.

A recent strand asks not only whether beliefs are distorted but whether the distortion covaries with position in the wealth distribution. Bardóczy and Guerreiro [2024] study macroeconomic stabilisation with imperfect expectations and show that the effectiveness of unemployment insurance depends on how beliefs about job loss are formed, which is the same margin our estimation measures. Our contribution here is to estimate the covariance directly: the belief wedge rises with liquidity, which identifies a gradient of the penalty in wealth, and we show that this gradient is separately identified from discount factor heterogeneity by the wealth distribution rather than by the distribution of marginal propensities to consume.

Our results on the fiscal multipliers also relate to Eichenbaum et al. [2026]: their survey evidence shows the planned propensity to spend out of transfers equals the MPC, meaning households do not price in future tax liabilities. They introduce a partial inattention framework in a model with uninsurable risk and get larger multipliers from inattention to future taxes.

There is a literature that applies behavioural frictions to the aggregate expectation operator. Gabaix [2020] and Angeletos and Huo [2021] modify how agents form expectations about aggregates, through cognitive discounting and through a learning representation of myopia and anchoring respectively. We use the representation of Angeletos and Huo [2021] as the machinery for our learning block, since it expresses an information friction as a transformation of the sequence-space Jacobian and is therefore agnostic about what generates that Jacobian. We apply those beliefs also to own income risk, beyond macro shocks. The representative agent comparison shows that the difference is quantitative and large: a distortion of aggregate beliefs cannot generate the asymmetry, because lifetime utility is orders of magnitude less sensitive to

the aggregate path than to own income risk.

The paper also relate to the survey-expectations literature. Our null results connect to Coibion et al. [2022], who find that households barely respond to central bank communication. Our evidence is consistent with theirs and extends it to a different object: not inflation expectations but beliefs about own labour market risk, which is the margin heterogeneous agent models actually use.

2 Estimating the Belief Process from Micro Data

The recent increase in availability fo survey data allows me to introduce a belief formation process that can be disciplined on micro data. We estimate the belief process by relying on data from the New York Fed Survey of Consumer Finance. This is one of the most reliable and best representative: it uses a rotating panel design that tracks the exact same individuals over time, rather than just taking random snapshots of different people each month. Below I describe in detail the specific questions I rely upon.

In terms of belief formation process, the estimation does not presuppose a specific belief model, but it identifies is a change of measure, and that object is common to essentially every theory of distorted beliefs. In other words, the estimation focuses on quantifying properties that are common to several belief formation processes.

Before described the detailed process, let us introduce some notation. Throughout, s denotes the household's exogenous idiosyncratic state and $e(s)$ the labour productivity associated with it. Hence $e(s)$ enters wage income while s indexes the state, histories are written $s^t = (s_1, \ldots, s_t)$. The symbol e standing alone is reserved for the exponential functions.

Let $\tilde{\pi}$ denote the household's subjective measure over next-period states and $\bar{\pi}$ the objective one, and assume $\tilde{\pi} \ll \bar{\pi}$, which implies that the household does not assign positive probability to events that cannot occur. Then there exists a Radon-Nikodym derivative m with $\bar{\mathbb{E}}[m] = 1$ such that $\tilde{\pi}(s') = \bar{\pi}(s')\, m(s')$, and writing $\psi \equiv \ln m$ every subjective measure can be represented as an exponential tilt:

$$\tilde{\pi}(s') \;=\; \frac{\bar{\pi}(s')\, e^{\psi(s')}}{\sum_{s''} \bar{\pi}(s'')\, e^{\psi(s'')}}, \qquad \text{so that} \qquad \underbrace{\ln \frac{\tilde{\pi}(s')}{\tilde{\pi}(s'')} - \ln \frac{\bar{\pi}(s')}{\bar{\pi}(s'')}}_{\text{log-odds wedge}} = \psi(s') - \psi(s''). \tag{1}$$

Given this general representation, specific theories of beliefs characterize ψ, and the theories divide into exactly two families depending on the degree to which households are responsive to the state. In most cases the responsiveness is proportional to the potential loss or through the value function attached by household to that state. In all those theories there is also some degree of anchoring to some fixed reference. The balance between the two identifies the theories: in the

estimation this criterion that identifies the two families is summarized by the two coefficients of 4. Further below, we explain how the belief process in our model captures those elements and is calibrated exactly on the estimated process.

- **Robust control and risk-sensitive preferences.** $\psi = \vartheta V$ [Hansen and Sargent, 2001, Tallarini, 2000]. The household over-weights a state in proportion its potential value-loss and chooses a posterior probability measure that minimizes entropy. There are typically two mathematical formulations, one with a multiplier and one with a constraint capturing ambiguity aversion: they deliver the same ψ when the set of priors is rectangular [Gilboa and Schmeidler, 1989, Epstein and Schneider, 2003].

- **Rational inattention.** ψ is proportional to the payoff and scaled by the shadow price of information. The agent again chooses a probability measure by minimizing entropy, where the prior enters multiplicatively. The optimal probability measure weights the value in each state against a common anchor, which is the average posterior across all states [Sims, 2003, Matějka and McKay, 2015]. The anchor produces a compression, which is indeed detectable in the data.[1]

- **Cognitive discounting.** Deviations from steady state are scaled by an attention parameter [Gabaix, 2020], so the default is the steady state and the shrinkage factor is the same in each state. The compression or shrinkage criterion is similar to the one in rational inattention except that it applies to the mean.

- **Sticky information.** A share of households act on a stale conditional, so the population belief is a mixture of the current objective measure and its own lag [Mankiw and Reis, 2002, Reis, 2006]. The default is the lagged belief rather than a fixed distribution, so the distortion is zero in a stationary environment and non-zero only along a transition.

- **Diagnostic expectations.** States are evaluated at the probability measure: $\psi = \gamma \ln \left[\bar{\pi}(s' \mid I_t)/\bar{\pi}(s' \mid I_{t-1})\right]$, the log likelihood ratio between the current and the previous conditional [Bordalo et al., 2018]. This is the mirror image of the previous item: sticky information

[1] Robust control and rational inattention both distort beliefs by minimising an objective containing a Kullback-Leibler term, and in both the solution is an exponential tilt with the payoff in the exponent. They differ in what the divergence is measured against, and that difference is what (4) detects. In the multiplier problem of Hansen and Sargent [2001] the penalty is: $\mathrm{KL}\big(\tilde{\pi} \,\|\, \bar{\pi}(\cdot \mid s)\big)$, measured against the true conditional for the state the household is in, so that conditional is the base of the tilt, $\tilde{\pi} \propto \bar{\pi}(\cdot \mid s)\, e^{\vartheta V}$. Subjective and objective log-odds then move one for one and the distortion is a pure level shift, $b = 1$. In rational inattention the constraint is mutual information, $I = \sum_s \mu(s)\, \mathrm{KL}\big(\tilde{\pi}(\cdot \mid s) \,\|\, \bar{q}\big)$ with $\bar{q} = \sum_s \mu(s)\tilde{\pi}(\cdot \mid s)$, so the reference is the household's own average conduct across states rather than the truth in any one of them [Sims, 2003, Matějka and McKay, 2015]. The base of the tilt is therefore the same $\bar{q}$ for every s, $\tilde{\pi}(\cdot \mid s) \propto \bar{q}\, e^{u(\cdot,s)/\lambda}$, movements in the true conditional are only partly passed through, and $b < 1$ by the weight on $\bar{q}$. The two are distinguished by whether whether the household's beliefs are anchored to the truth or to itself.

pulls the belief toward the lag, diagnostic expectations push it away, over-weighting states that have just become more likely. This too operates only in the dynamic.

Two implications matter for what follows. First, the log-odds wedge on the left of (1) is estimable without having to specialize exactly to any of the above theories, since it requires only a subjective and an objective probability for the same pair of states and the same household. In summary, the estimation delivers a process for the beliefs that quantifies the combination of the features common to the above theories.

Let us briefly describe what criteria guide the generalized beliefs specification adopted in the macro model. First, the belief distortion is operative both in the steady state and along the dynamic: it seems very plausible to argue for consistency in limited capacity; second the steady state of a model with uninsurable risk assumes that the agent has to solve complex forecast problems (Moll [2026]), and limited capacity, affecting also precautionary behaviour, appears a sensible assumption. The process introduced in the model is calibrated on the one estimated below. The sign of the belief wedge, its persistence, its cross-sectional gradient and its behaviour around identified shocks are properties of (1) alone. The functional form specializes two characteristics. First, $\psi = \vartheta V$ and therefore the tilt depends on households' value function; a different theory would map the same wedge into a different parameter. Second, it models a compression, calibrated in the data, that is meant to capture the convergence (common to most of those theories) toward a common posterior.

Data. For the estimation we use data from the core module of the New York Fed Survey of Consumer Expectations runs monthly from June 2013. The three public vintages give 176,101 household-months on 23,369 households through December 2024. Households enter a rotating panel and are interviewed for up to twelve consecutive months, which is what makes it possible to observe both a household's stated probability and its own subsequent realisation.

In the model the earning risk are modeled to capture both variation in the level, but also in the employment status of the household (see Krusell and Smith [1998]. Expectations in the model are formed about both the idiosyncratic state (hence they are operative in the steady state) and about the aggregate state, through a dynamic learning component. Therefore for the process estimation we leverage also on questions about employment prospects and status, which form our outcome variables (and in some cases regressors as specified below). We use four items for the estimation. *Q13new* is the percent chance of losing one's main job over the next twelve months, hence captures expectations about unemploymet risk. *Q22new* is the percent chance of finding an acceptable job within three months conditional on losing the current one. *Q10_1* through *Q10_10* record current employment status, from which realised transitions are constructed in the data and in the model. `Q24_bin1` through `Q24_bin10` give the subjective

density over own earnings growth conditional on staying in the job.

We use liquidity as a sorting variable, so the wedge can be computed by liquidity quintile and we get cross-sectional variation. Liquidity is also one of the controls in the first stage that predicts realised separation from observables. This variable is taken from the Credit Access module, fielded every four months. Item *N25* is the percent chance the household could raise $2,000 for an unexpected need within a month. This is preferable to a wealth level because the model's cross-sectional prediction concerns proximity to the borrowing constraint, which *N25* measures directly. Furthermore, no clean liquid wealth level exists in the public files. Households are sorted on their time-averaged value so that the sort is predetermined.

The Household Spending module, fielded every four months from December 2014, gives realised spending growth and an elicited marginal propensity to consume out of a ten percent income gain. Both are used in one of the estimation steps to measure the extent of the compression, or convergence toward a common distorted posteriors, which (as argued above) is common to many of the theories.

Estimation Specification for the Main Belief Process. The generalized beliefs' functional form that we adopt in the model features an exponential tilt of the objective transition probabilities of the following form:

$$\tilde{\pi}(s' \mid s, a) \;=\; \frac{\bar{\pi}(s' \mid s)\, e^{\vartheta V(s',a)}}{\sum_{s''} \bar{\pi}(s'' \mid s)\, e^{\vartheta V(s'',a)}}. \tag{2}$$

The normalising constant cancels from any ratio of two states, so:

$$\underbrace{\ln \frac{\tilde{\pi}(s')}{\tilde{\pi}(s'')} - \ln \frac{\bar{\pi}(s')}{\bar{\pi}(s'')}}_{\text{log-odds wedge}} \;=\; \vartheta \left[V(s', a) - V(s'', a) \right]. \tag{3}$$

Equation (3) is linear, requires no assumption beyond the functional form in (7), and returns ϑ in the units the model uses. We take s' to be job separation and s'' to be continued employment. The subjective expectation are measures using *Q13new*. Because the observed panel window h_i varies across households, the twelve-month probability is rescaled under a constant monthly hazard, $\tilde{p}_i(h_i) = 1 - (1 - \tilde{p}_i^{12})^{h_i/12}$, and the sample is restricted to households with at least three months of lookahead. The objective probabilities are measured with the household's own realised transition within the same panel. This is important as it allows me to maintain measure from the same panel, as opposed to match with administrative measures.[2] Standard errors are

[2]The wedge is a difference of log-odds. The survey does not supply a probability but one draw from it, so $\bar{p}_i \in \{0, 1\}$ and $\text{logit}(\bar{p}_i) = \pm\infty$. Therefore, the objective probability must be estimated as a frequency. Households are pooled into cells c formed by liquidity quintile crossed with month and the wedge is:

$$\omega_c \;=\; \text{logit}\Big(\bar{\tilde{p}}_c\Big) \;-\; \text{logit}\Big(\widehat{\text{freq}}_c\Big),$$

the mean stated probability against the share of the cell who actually separated. The pooling also allows to handle the zeros in the responses. The expression takes the logit of a mean rather than the mean of logits, and by Jensen's inequality these differ. Since we cannot include household fixed effects to a cell-level object, we verify the

from a household-clustered bootstrap with 400 replications.

The estimation sample is 69,806 observations on 11,233 households, June 2013 to September 2024.

The wedge in the data. Table 1 reports the two probabilities and their log-odds wedge, by liquidity quintile.

Liquidity	$\tilde{p}$	$\bar{p}$	Wedge	(s.e.)	ΔV
Q1 (least liquid)	0.1002	0.0780	0.275	(0.059)	3.35
Q2	0.1014	0.0577	0.611	(0.069)	2.65
Q3	0.0931	0.0396	0.911	(0.082)	2.06
Q4	0.0762	0.0245	1.187	(0.105)	1.53
Q5 (most liquid)	0.0666	0.0260	0.984	(0.115)	0.93

Table 1

Subjective and realised twelve-month job separation probabilities and the log-odds wedge, by liquidity quintile. ΔV is the model value difference between employment and non-employment at the matching wealth percentile.

The subjective probability exceeds the realised frequency in every quintile: households are pessimistic about their own employment risk, which is $\vartheta < 0$, the sign the model requires. Every estimate is significant, with t statistics between 4.5 and 11.3.

The sign is confirmed independently on the job-finding margin. Comparing the stated *Q22new* of households, who subsequently separate, against their own realised re-employment three months later, on a flow sample of 1,770, gives a subjective 0.5418 against a realised 0.5475, a log-odds wedge of -0.024. This is mild pessimism about finding work, which is what a common $\vartheta < 0$ implies for the reverse transition.

We note however two things. The wedge rises with liquidity while ΔV falls, so the ratio $\hat{\vartheta} = \omega_c / \Delta V_c$ implied quintile by quintile ranges from -0.08 to -1.06, an order of magnitude apart and far outside the standard errors. And the base rate $\bar{p}$ falls with liquidity alongside ΔV, from 0.078 to 0.026. This suggests that the wedge may contain other components, and since the goal is to use those parameters to calibrate the belief formation process of the model exactly, we shall purge those components. This is done in the next section.

Separating the distortion from true risk. The wedge, formally $\vartheta \Delta V$ in the specification, may also contain components of the subjective probability that are unrelated to the household's actual risk. Those may include reference anchoring for instance. Separating those component is crucial for the model's calibration. We separate those components in three steps.

First, realised separation is projected on age, education, income band, region, liquidity

estimates are stable as cells are enlarged from about 100 to about 3,000 observations. Where a household-level measure is needed, in the behavioural test below, we use a different object, the level difference $\tilde{p}_i - \hat{p}_i^{\text{true}}$ between the stated probability and the risk predicted from observables, which is defined for every household.

quintile, household size, panel tenure and the observed horizon. The fitted value $\hat{p}_i^{\text{true}}$ is the household's predicted separation risk given observables.

Second, under (7) the subjective and objective log-odds must move one for one with the tilt entering as a level shift, hence for identification the coefficient b in:

$$\text{logit}(\tilde{p}) \;=\; a \;+\; b\,\text{logit}(\bar{p}) \;+\; c\,\Delta V \tag{4}$$

should equal one, and $\hat{\vartheta} = -c$. Estimated across liquidity-by-year cells, $b = 0.109$ with a standard error of 0.047. Therefore is definitely rejects the restriction to one. Subjective beliefs are compressed: they respond far less than one for one to true probabilities. This is the component of the beliefs that links the belief formation process to the anchoring components of the theories spelled out above (see Bordalo et al. [2018], Gabaix [2020], Sims [2003]). The result is robust even when survey weights are applied and when all focal responses are dropped, and it is stable as cells are enlarged from about 100 to about 3,000 observations, so it is neither due to weighting, nor rounding, nor attenuation bias. Including $\text{logit}(\bar{p})$ is what stabilises the estimate: it absorbs the compression, which scales with the base rate rather than with ΔV, and which would otherwise load onto ΔV because the two fall together across the quintiles of Table 1. Conditioning on it, $\hat{\vartheta}$ lies between -0.11 and -0.19 across every specification, with standard errors near 0.05, in place of the -0.08 to -1.06 implied by the raw wedge.

Third, the estimate $b = 0.109$ establishes that stated beliefs are compressed, but not whether this compression actually affects behaviour. We separate the two by testing whether the stated probability predicts behaviour once the risk the household actually faces is controlled for. Once we control for liquidity, realised spending growth loads on predicted true risk at $t = 7.6$ against $t = 6.2$ on the stated belief, expected spending growth at $t = 10.3$ against $t = 0.4$, and the elicited marginal propensity to consume on neither. Behaviour tracks states of nature rather than stated beliefs, so the bulk of the measured compression is survey response behaviour; consistent with this, the individual-level slope of the stated probability on predicted own risk is 0.771, implying a compression of 0.229 rather than 0.89. We therefore retain the smaller structural component, $\kappa_c = 0.23$, which we use to calibrate the model. Note that those parameters affect precautionary saving and the MPC, hence their exact calibration is quite consequential.

The survey therefore delivers four of the five belief parameters: the level ϑ from the coefficient c, the asset gradient γ from the cross-sectional profile of the wedge, the persistence $\rho_\vartheta = 0.787$ from the autocorrelation of the monthly series, and the compression $\kappa_c = 0.23$. The process in the model will also feature a dynamic component, the parameter on that is calibrated to match main facts on the macro responses.

3 The model

The model is an economy in which households face uninsurable risk and borrowing constraints, invest in one asset and choose consumption-saving and labor supply. The main departure from the benchmark economies lies in the introduction of the belif process. We describe this part in detail. The rest of the model, which we summarize briefly, is a standard New Keynesian model.

Households. Households' preferences are separable: $u(c,n) = \frac{c^{1-1/\sigma}}{1-1/\sigma} - \varphi\frac{n^{1+1/\nu}}{1+1/\nu}$. A household in idiosyncratic state s with beginning-of-period assets a_- faces the following budget constraint:

$$c_t + a_t \;=\; (1+r_t)\,a_{t-1} + w_t\,e(s_t)\,n_t + T_t(s_t), \qquad a_t \geq 0, \tag{5}$$

where $e(s)$ is the labour productivity associated with state s and $T_t(s)$ collects dividends received and taxes paid, both distributed in proportion to productivity. The state s follows a Markov chain with objective transition $\bar{\pi}(s' \mid s)$.

The household chooses an allocation and a subjective measure. Writing the problem recursively makes the two choices sequential within the period (see Bassanin et al. [2021]). Given the continuation value, the household first selects the measure it will use to evaluate the future, and then chooses consumption, labour supply and saving under that measure. Note that this is a measure of the earning idiosyncratic states and as such its distortion operates in the stochastic steady state. Formally:

$$V(s,a_-) \;=\; \max_{c,\,n,\,a}\Big\{\,u(c,n) \;+\; \beta \underbrace{\min_{\pi(\cdot|s)}\Big[\textstyle\sum_{s'}\pi(s'\mid s)\,V(s',a) \;+\; \frac{1}{|\vartheta|}\,\mathcal{K}\big(\pi(\cdot\mid s)\,\big\|\,\bar{\pi}(\cdot\mid s)\big)\Big]}_{\text{stage 1: optimal beliefs}}\Big\} \tag{6}$$

subject to (5), where $\mathcal{K}$ is the Kullback-Leibler divergence and $\frac{1}{|\vartheta|}$ prices deviations from the objective measure. Appendix A shows that (6) is equivalent to the sequence problem, that the operator is a contraction, and that the resulting plan is dynamically consistent.

Stage 1: optimal beliefs. The inner problem is static and convex. Minimising over $\pi(\cdot \mid s)$ subject to $\sum_{s'}\pi(s' \mid s) = 1$ yields the exponential tilt:

$$\tilde{\pi}(s' \mid s,a) \;=\; \frac{\bar{\pi}(s' \mid s)\,e^{\vartheta V(s',a)}}{\sum_{s''}\bar{\pi}(s'' \mid s)\,e^{\vartheta V(s'',a)}}, \qquad \vartheta \equiv -\tfrac{1}{|\vartheta|} < 0, \tag{7}$$

and substituting it back collapses the bracket into a risk-sensitive certainty equivalent, so that:

$$V(s,a_-) = \max_{c,n,a}\Big\{u(c,n) + \frac{\beta}{\vartheta}\ln\bar{\mathbb{E}}\Big[e^{\vartheta V(s',a)}\,\Big|\,s\Big]\Big\}. \tag{8}$$

Note that the asset choice a enters the tilt, since the household evaluates next period's states at the wealth brought to next period. The two stages are therefore not independent, and (6) must be solved jointly rather than sequentially in the numerical implementation.

Stage 2: consumption, labour supply and saving. Given the tilted measure, the household solves a standard problem. Let $\mu_t \geq 0$ be the multiplier on the borrowing constraint. The first-order conditions are an intratemporal condition equating the marginal rate of substitution to the effective wage:

$$\varphi\, n_t^{1/\nu} \;=\; w_t\, e(s_t)\, c_t^{-1/\sigma}, \tag{9}$$

and an Euler equation in which the tilt appears as a wedge on the stochastic discount factor:

$$u_c(c_t, n_t) - \mu_t = \beta\, \bar{\mathbb{E}}_t\big[m_{t+1}\, u_c(c_{t+1}, n_{t+1})(1 + r_{t+1})\big], \qquad m_{t+1} = \frac{e^{\vartheta V(s_{t+1}, a_t)}}{\bar{\mathbb{E}}_t\, e^{\vartheta V(s_{t+1}, a_t)}}, \tag{10}$$

with complementary slackness $\mu_t a_t = 0$. The envelope condition is $V_a(s_t, a_{t-1}) = (1+r_t)u_c(c_t, n_t)$, so it is the marginal value of wealth that the tilt reweights. Labour supply (9) is static and therefore carries no belief distortion directly. Beliefs affect hours only through c_t. Setting $\vartheta = 0$ gives $m_{t+1} = 1$ and recovers the rational expectations problem exactly.

Lemma 1 (Direction of the Belief Wedge; Bassanin et al. [2021])**.** *If $\vartheta < 0$ then $m_{t+1} > 1$ for low-value states and $m_{t+1} < 1$ for high-value states: beliefs are left-skewed and households act with pessimism. At $\vartheta = 0$ the tilt is degenerate and beliefs are rational.*

Equation (7) is an exponential tilting of the reference measure, also known as the Esscher transform. Since $\bar{\mathbb{E}}_t[m_{t+1}] = 1$ by construction, m_{t+1} is the Radon-Nikodym derivative of the subjective with respect to the objective measure, and in the terminology of Hansen and Sargent [2001] a martingale distortion. Note that the tilt parameter is $\vartheta < 0$: this is the penalty on deviations from the objective measure is $\theta = -1/\vartheta > 0$, or the multiplier, and is what appears in (6). In Appendix A we present the recursive formulation of the households' optimization problem: the multiplier on the discounted constraints there is $(1-\beta)/|\vartheta|$, positive, since written with the opposite sign the penalty would reward deviations and the inner minimisation would be unbounded.

Anchoring. The estimation identifies a component of the subjective measure that is insensitive to the household's own state. We allow for it as a mixture with the stationary distribution:

$$\hat{\pi}(s' \mid s) = (1 - \kappa_c)\, \bar{\pi}(s' \mid s) + \kappa_c\, p(s'), \tag{11}$$

which replaces $\bar{\pi}$ in (7) and in the penalty of (6). At $\kappa_c = 0$ the model is unchanged. Corollary 1 of Appendix A verifies that the mixture leaves the recursive structure and the contraction property intact.

Dynamic Wedge: Learning over Income-Risk Regimes. The subjective probability described so far is a measure of the earning states. However, our evidence pointed out at a wedge

component that does not directly relate to the individual household state. We therefore introduce a second component of the belief formation process which is operative over the dynamic, and acts as a Bayesian learning in which the tilt distorts the standard update formula. As we shall see this component also allows the model to match some patterns of the impulse response and the transmission of shocks.

We define the state space of this learning update as the number of regimes in the risk-space, the volatility of the earning process.[3] Let the economy be in one of two income-risk regimes, and let households learn which one prevails. Write p_τ for the objective posterior τ periods after the shock, growing at rate λ. The robust household tilts that posterior with the same ϑ:

$$\tilde{p}_\tau = \frac{p_\tau e^{\kappa}}{p_\tau e^{\kappa} + (1 - p_\tau)}, \qquad \kappa = \vartheta \, \Delta V^{\text{regime}}, \tag{12}$$

where ΔV^{regime} is the household's valuation of the regime difference. Note that $\tilde{p}_\tau = 1$ whenever $p_\tau = 1$: a degenerate posterior cannot be tilted, which is the perfect-foresight recovered as a limiting case rather than imposed as a prohibition. The functional form of the tilt is equivalent to the one considered in the long run beliaf formation with respect to the idiosyncratic earning states. Therefore the belief formations process are present in the long and the short run and are consistent to each other. What distinguishes them is that the one that applies at short horizons captures tilts on the update.

Learning and Dynamic Responses. One can show anlytically how the learning process may affect the dynamic response to shocks in our environment. Setting $E[\tau, s] = \tilde{p}_\tau$ and applying the Angeletos-Huo representation [Angeletos and Huo, 2021] transforms the full-information Jacobian M into:

$$M^{\text{beh}}_{t,s} = \sum_{\tau=0}^{\min(t,s)} \big(E[\tau, s] - E[\tau - 1, s]\big) \, M_{t-\tau,\, s-\tau}. \tag{13}$$

The term $M_{t-\tau,s-\tau}$ is the response of a household that learned about the date-s shock exactly τ periods late. This is the only component of the model that operates in calendar time and it is largely responsible for shaping the dynamic response of the economy. This is why we will validate this component by showing that it match certain regularities of the dynamic response, which would otherwise be unmatched. Because κ flips sign with the sign of the shock, the friction is asymmetric: for an easing $\Delta V^{\text{regime}} > 0$ so $\kappa < 0$ and the pessimist underweights good news, while for a tightening the pessimist overweights bad news.

[3] Models with uninsurable risk are generally solved forward, at the aggregate level, under perfect foresight or with an approximation around an MIT shock. This implies that formally there is not uncertainty or probability measure over the dynamic. We therefore engineer the states with the regime shifts across different levels of risk. Households then attach Bayesian updates to those states, and the updates are distorted by the tilt.

3.1 The rest of the model

The rest of the model is standard. Production is by a competitive final good producer aggregating varieties with elasticity $\mu/(\mu-1)$, each produced with labour, $L_t = Y_t/Z_t$, subject to a Rotemberg price adjustment cost, giving:

$$\ln(1+\pi_t) = \kappa_p\left(\frac{w_t}{Z_t} - \frac{1}{\mu}\right) + \frac{Y_{t+1}}{Y_t}\frac{\ln(1+\pi_{t+1})}{1+r_{t+1}}. \tag{14}$$

Dividends are output net of the wage bill and the adjustment cost. Government debt is constant at B with a balanced budget, $T_t = r_t B$. Monetary policy follows $1+r_t = (1+r^*_{t-1}+\phi\pi_{t-1})/(1+\pi_t)$. The distribution $D_t(s,a)$ evolves under the *objective* transition matrix and the saving policy, and markets clear, $A_t = B$, $N_t = L_t$. In Section 4.1 I consider an extension that adds capital and variable utilisation. Furthermore, we consider an extension with time-varying public debt and a debt rule to assess also the transmission mechanism of fiscal shocks and the fiscal multipliers.

3.2 Solution Method and Calibration

The stochastic steady state is solved with the endogenous gridpoint method of Carroll [2006] and the dynamic response to shock is solved with the sequence-space Jacobian method of Auclert et al. [2021]. Implementing the tilt inside that framework requires some transformation. The toolkit takes the objective expectation over next-period states before the household's the envelope (backward step in the language of this method) is evaluated, and the tilt requires those states to be reweighted individually by $e^{\vartheta V(s',a)}$ before the sum is taken. Appendix B sets out the transformation that resolves it. With this transformation no modification of the toolkit is required, and the distribution continues to evolve under the objective transition while beliefs are tilted.

Calibration. The income process is calibrated to the same panel used for the belief estimation. An explicit non-employment state replaces the practice of reading the bottom point of a discretised productivity grid as job loss, which implied a twelve-month entry probability of 0.0062 against 0.083 in the data, a factor of thirteen.[4] The paraneter ς indicates the transition from employment to non-employment, the parameter f indicates transitions from non-employment to employment.[5]

Nest, we calibrate the beliefs' parameters, based on our estimates. Three of the four belief parameters are estimated, and the fourth, which controls the dynamic learning component, is inferred from the model's ability to match aggregate responses. From (4), estimated after

[4] States 1 to 7 are employed productivity from a Rouwenhorst discretisation; state 0 is non-employment with income ρ_r times mean employed earnings.

[5] The Rouwenhorst matrix scaled by $1-\varsigma$ within employment and re-entry drawn from the stationary employed distribution.

Parameter	Target	Data	Model
ς	quarterly separation rate	0.0215	0.0215
	twelve-month separation	0.0833	0.0833
f	quarterly job finding	0.2992	0.2992
	non-employment share	0.0583	0.0670
ρ_r	replacement rate	not targeted	0.517

Table 2

Income process, calibrated to SCE panel transitions. The non-employment share is implied rather than targeted; $\varsigma/(\varsigma + f) = 0.067$ against an observed 0.058.

controlling for the base rate, $\hat{\vartheta}$ lies between -0.11 and -0.19 across every specification; the preferred estimate, weighted and excluding focal responses, is -0.188 with a standard error of 0.053. We use -0.19. The first-order autocorrelation of the monthly $\hat{\vartheta}_t$ series is 0.787; a regression of $\hat{\vartheta}_t$ on its own lag and payroll growth gives 0.755. The implied penalty is larger in absolute value for more liquid households. Expressed through the model's asset profile $\vartheta_t(a) = \vartheta_t + \gamma(x(a) - \frac{1}{2})$ with $x(a) = a/(a + a_{\text{ref}})$, the estimated spread across the wealth distribution is about 0.066, so $\gamma \approx -0.066$: richer households are more cautious, after controlling for risk. Finally, equation (12) defines $\kappa = \vartheta \Delta V^{\text{regime}}$, and ΔV^{regime}, which is the tilt applied to posteriors across risk regimes. It is computed from the model solved ss the change in the distribution-weighted value function across income-risk regimes:

Regime	ΔV^{regime}	$\kappa = \vartheta \Delta V$
σ_e +10%	-0.570	0.108
σ_e +20%	-1.108	0.211
σ_e +30%	-1.619	0.308

We use the twenty percent regime, $\kappa = 0.21$. The individual-level slope of stated probability on predicted own risk is 0.771, implying compression of 0.229. The cell-level regression suggests 0.89, but the behavioural test in Section 2 shows that the bulk of that is survey response behaviour rather than a belief primitive, so we retain the smaller structural value. At last. the rate at which the regime posterior converges, λ, corresponds to an average of four quarters to learn the regime.

All remaining parameters are from Auclert et al. [2021]: $r = 0.005$ quarterly, $\sigma = 0.5$, $\nu = 0.5$, $B = 5.6$, $\mu = 1.2$, $\kappa_p = 0.1$, $\phi = 1.5$, $\rho_e = 0.966$, $\sigma_e = 0.5$. For each belief configuration β and φ are re-solved so that the asset market clears at $A = B$ and effective labour supply equals one. The tilt shifts precautionary saving substantially, so comparisons at a fixed β would confound the belief mechanism with a different steady-state wealth distribution.

4 Results

We start by presenting the comparison between the baseline heterogenous agent model with one asset and the one augmented with the estimated beliefs. We rely primarily on impulse response functions as they can outline some stylized patterns that make the model comparable to the data for validation.

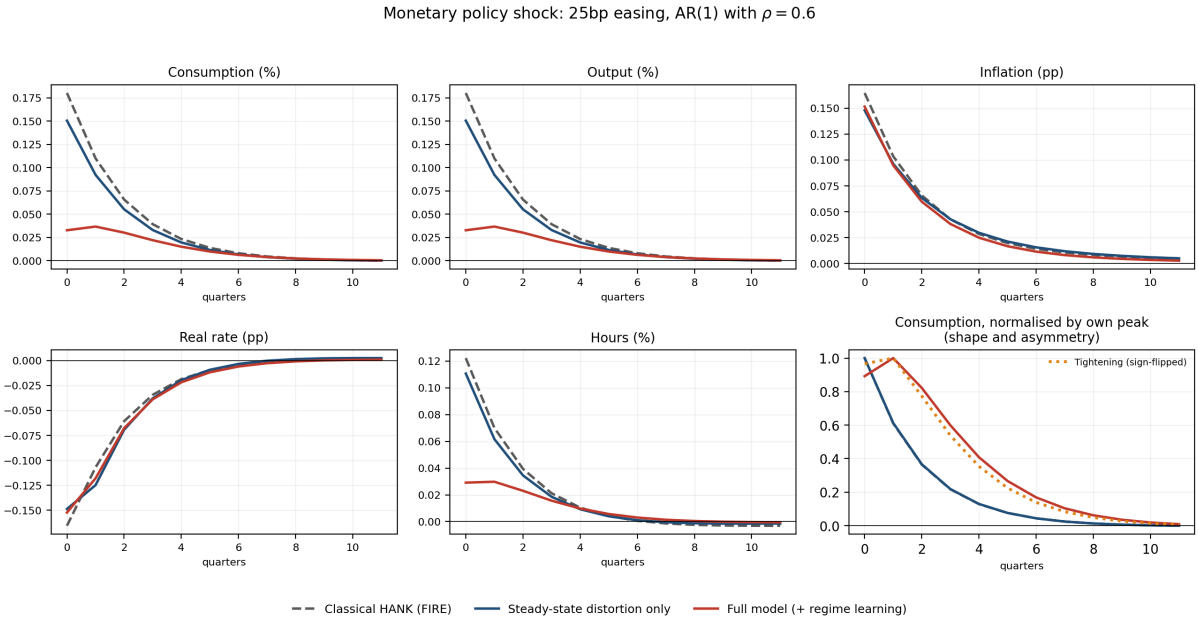


Figure 1

Impulse responses to a 25 basis point monetary easing with quarterly persistence 0.6. The bottom right panel normalises each consumption response by its own peak to display the shape, and adds the sign-flipped response to a tightening of identical size.

	β	Impact (%)	Peak (%), quarter	Cumulative 12q
Classical HANK (FIRE)	0.98017	+0.1801	+0.1801, q0	+0.4475
Steady-state distortion only	0.97807	+0.1504	+0.1504, q0	+0.3749
Full model, with regime learning	0.97807	+0.0326	+0.0365, **q1**	+0.1602
Tightening, sign-flipped	0.97807	+0.0446		+0.1925

Table 3

Consumption response to a 25 basis point easing. Marginal propensity to consume: 0.096 under FIRE, 0.165 under the estimated distortion.

Figure 1 shows impulse response functions to a 25 basis point monetary shock. The baseline heterogenous agent model is the dashed line. Consumption jumps 0.18 percent on impact and decays monotonically, and the peak is at impact. This is the shape that Christiano et al. [2005] and the subsequent literature identified as the failure of this class of models against the data, which shows dampened and hump-shaped responses. The estimated belief model differs in three respects.

Once the model is augmented with the estimated belief process the impact response in

consumption falls from 0.180 to 0.033 percent. Consumption peaks in quarter one rather than quarter zero. The hump is primarily driven by the Bayesian learning process, so by the part of the distorted beliefs that moves in calendar time. A remarkable feature is that this comes without any additional assumption on discount factors, asset return processes or preferences' assumption. The steady-state distortion alone accounts for only a sixth of that, taking the response to 0.150; the bulk comes from learning, because most households have not yet revised their view of the risk regime in the quarter the shock lands. The cumulative response falls from 0.448 to 0.160.

Another remarkable feature of the impulse responses that match empirical evidence on the transmission mechanism of policy shocks is the asymmetry of the response. A contraction of identical size produces a cumulative response of 0.193 against 0.160 for an easing, a ratio of 1.20, and peaks earlier. Because $\kappa = \vartheta \Delta V^{\text{regime}}$ flips sign with the shock, a pessimist underweights good news and overweights bad news, so contractions transmit faster and more strongly. No symmetric friction, including sticky information, cognitive discounting or level-k thinking, produces this. It is consistent with the empirical literature on asymmetric monetary transmission [Tenreyro and Thwaites, 2016].

Reconciling Micro Elasticities and Macro Dynamic. A key discipline devise of this class of model lies in verifying whether the model can match at the same time micro elasticities, alongside macro dynamics. On the micro side, the marginal propensity to consume rises from 0.096 under rational expectations to 0.165 under the estimated distortion, placing it inside the 0.15 to 0.20 range of the recent experimental estimates. The level of the household MPC is actively disputed: Boehm et al. [2025] or Colarieti et al. [2024] report a one-month MPC of 0.23 0.23 in a randomized experiment, while Orchard et al. [2025] argue that the leading micro estimates are upward biased and correct them downward. Our quarterly 0.165 lies between these, though the comparison with the monthly figure of Boehm et al. [2025] is not like for like. Nothing was targeted at the marginal propensity to consume, so the movement rather than the level is the informative object. Remarkably, nothing was targeted at the marginal propensity to consume, so this is an over-identifying check rather than a fitted moment. The model therefore raises the micro elasticity while lowering the macro elasticity.

4.1 Extensing with Capital and variable utilisation

The model of Section 3 has no accumulating state on the production side. It is useful to verify that the dynamic and the transmission uncovered so far remain valid under this extension, which has now become common ground for heterogenous agents models. Furthermore, its presence, alongside with other persistence inducing devices, such as adjustment costs and variable uti-

lization of capital may help to get the macro dynamic even closer to the empirically relevant one.

We therefore extend the model as follows. Capital is held inside the firm, financed out of retained earnings, and households receive the residual as the lump-sum transfer D_t that already appears in their budget constraint.[6] We describe below only the new parts of the model and calibration.

Output uses effective capital $u_t K_{t-1}$,

$$Y_t = Z_t L_t^{1-\nu} \left(u_t K_{t-1}\right)^{\nu}, \tag{15}$$

with capital accumulation subject to a convex adjustment cost governed by: ε_I,

$$Q_t = 1 + \frac{1}{\delta \varepsilon_I}\left(\frac{K_t}{K_{t-1}} - 1\right), \qquad I_t = K_t - (1-\delta)K_{t-1} + \frac{K_{t-1}}{2\delta\varepsilon_I}\left(\frac{K_t}{K_{t-1}} - 1\right)^2. \tag{16}$$

Utilisation is chosen to equate the marginal product of effective capital to the marginal utilisation cost, with $a(1) = 0$ and $a'(1)$ equal to the steady-state rental rate:

$$\nu \, mc_t \frac{Y_t}{u_t K_{t-1}} \;=\; a'(u_t) \;=\; r^k_{ss}\, e^{\psi_u (u_t - 1)}, \tag{17}$$

and the resource cost $a(u_t)K_{t-1}$ enters the goods market clearing condition.

Calibration. ε_I is set so that the model reproduces the relative size of the investment and consumption responses to a monetary shock, which Christiano et al. [2005] put at roughly three to one. We use $\varepsilon_I = 4$; for a 25 basis point shock the ratio to consumption is the object the literature pins down and it is in the right range. We set $\psi_u = 1$.

ε_I	$C(0)$	$I(0)$	ratio
0.5	+0.1663	+0.102	0.61
1.0	+0.1636	+0.195	1.19
2.0	+0.1583	+0.355	2.24
4.0	+0.1494	+0.597	**3.99**
8.0	+0.1372	+0.890	6.48

	$C(0)$	Peak, quarter	Cumulative	$I(0)$
No capital, full information	+0.1801	+0.1801, q0	+0.4475	
No capital, estimated beliefs	+0.1504	+0.1504, q0	+0.3749	
+ capital	+0.0953	+0.0953, q0	+0.3511	+0.732
+ capital + utilisation	+0.1160	+0.1160, q0	+0.3998	+0.852
Full, + regime learning	+0.0246	+0.0288, **q1**	+0.1920	+0.871

Table 4

Consumption response to a 25 basis point easing across models. $\vartheta = -0.19$, $\kappa_c = 0.23$, $\varepsilon_I = 4$, $\psi_u = 1$, $\lambda = 0.75$.

Capital and investment do not seem to add much in terms of persistence. One conjecture

[6] Like before, asset market clearing is given by $A_t = B$. This similarity makes the comparison across the rows of Table 4 possible.

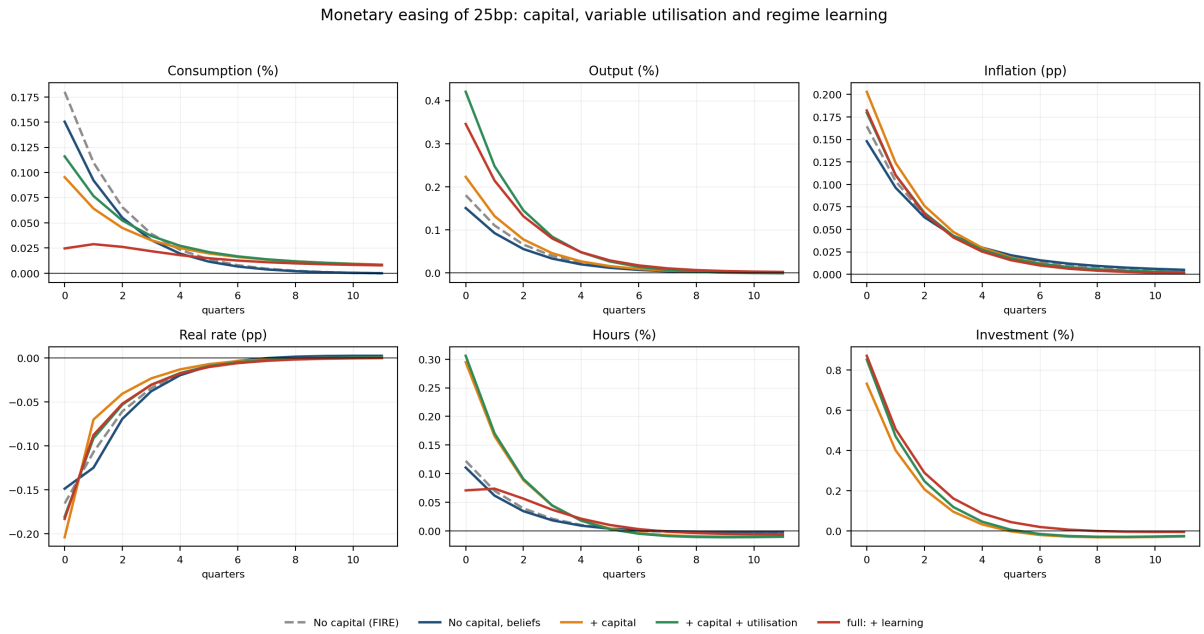


Figure 2

Monetary easing of 25 basis points. Each line reports the results for a different model. The household side is identical across all rows.

was that the dividends rebated to households would carry some of the persistence inherited from the firms' policy functions. That effect seems to be small.

Adding capital alone cuts the impact consumption response from 0.150 to 0.095, because investment now competes for resources. Hence capital dampens. Utilisation, however, restores part of it, to 0.116, by letting the firm raise effective capital without paying the adjustment cost, which flattens the marginal cost response. The two work in opposite directions on the level and neither touches the shape.

In sum, hump shaped and persistence are mainly driven by learning and are unaffected by capital.

5 Comparing with and without Uninsurable Risk

The model combines two ingredients, distorted beliefs and uninsurable idiosyncratic risk, and it is natural to ask which of those features is responsible for each of the model's properties we observe. We therefore run the same belief process in a representative agent version of the model. The nominal block, the policy rule and the steady-state targets are unchanged; only the household side is replaced, with $\beta = 1/(1+r)$ so that the two economies share the same steady state.

Let us first discuss how this changes the transmission mechanism and the channels that operate in the model. First, the steady-state tilt (7) is inoperative in a representative agent

model: the steady state beliefs are function of uninsurable risk. Absent that there is no stochastic state to form belief upon. However, the dynamic belief distortion, ΔV^{regime}, remains in (12) and must be computed from the aggregate path rather than from income risk.

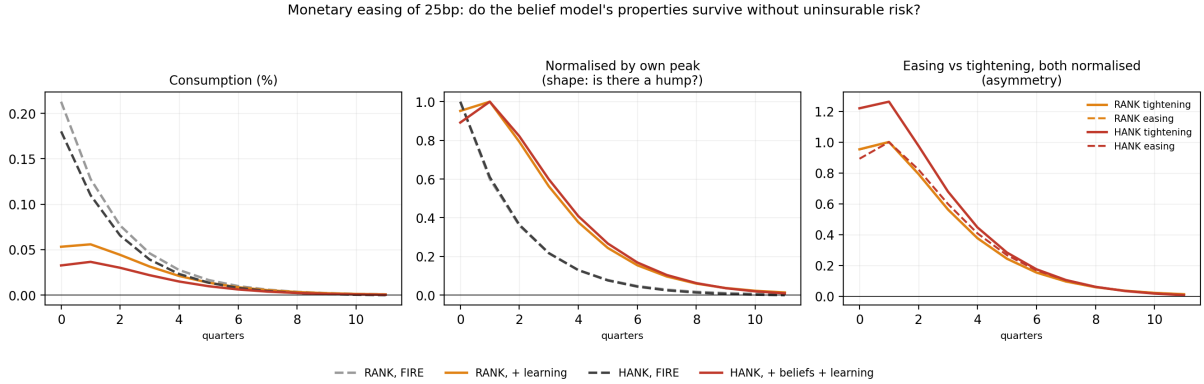


Figure 3

Monetary easing of 25 basis points. Left: consumption. Centre: each response normalised by its own peak, to display shape. Right: easing and tightening of identical size, both normalised by the easing peak, to display asymmetry.

	$C(0)$	Peak, quarter	Cumulative	Tightening / easing
RANK, full information	+0.2128	q0	+0.5308	
RANK, + learning	+0.0532	**q1**	+0.2406	1.001
HANK, full information	+0.1801	q0	+0.4475	
HANK, + beliefs + learning	+0.0326	**q1**	+0.1602	**1.201**

Table 5

The same belief process in a representative agent and a heterogeneous agent model. The final column is the ratio of the cumulative response to a contraction to that of an easing of identical size.

The representative agent model with learning peaks in quarter one exactly as the heterogeneous agent model does, and the dampening is of similar magnitude, from 0.213 to 0.053 against 0.180 to 0.033. In other words learning, the dynamic component of the distortion, is responsible for the hump and it is so both in the heterogenous and the representative agent model.

On the other side, the ratio of contraction to expansion is 1.001 in the representative agent model against 1.201 in the heterogeneous agent model. In other words, only the combination of belief distortion and uninsurable risk delivers the asymmetric response to policy shocks. Asymmetry needs $\kappa = \vartheta \, \Delta V^{\text{regime}}$ to be non-negligible, and the two models have entirely different objects available to put in ΔV^{regime}. Valuing the aggregate consumption path of a 25 basis point easing at the representative household's marginal utility gives 3.7×10^{-3} utils, against 1.108 utils for the twenty percent rise in idiosyncratic income risk used in Section 3.2. The implied penalty therefore falls from -0.210 to -7.0×10^{-4}, a factor of roughly three hundred, and the tilt on the posterior is inactive.

5.0.1 Beliefs and the Fiscal Multiplier

The overarching goal of the paper is to evaluate the role of distorted beliefs, alongside uninsurable risk, to affect the policy transmission mechanism. We do not limit ourselves to the monetary transmission mechanism, but explore also the consequences for the fiscal multiplier. To this purpose we modify the model to include a more meaningful fiscal sector. The baseline model holds government debt fixed and balances the budget period by period, which is a strong assumption for a fiscal experiment. We therefore compare it with a rule under which spending is debt-financed on impact and taxes adjust gradually:

$$T_t = T_{ss} + \phi_B\big(B_{t-1} - B_{ss}\big), \qquad B_t = (1 + r_t)B_{t-1} + G_t - T_t, \tag{18}$$

with $\phi_B = 0.10$ per quarter. The rule is passive in the sense of Leeper [1991], so determinacy is unaffected, and $\phi_B > r$ guarantees that debt converges (see also Auclert et al. [2018]). Steady-state government spending is twenty percent of output and the two regimes share the same steady state, so the comparison isolates the financing margin.

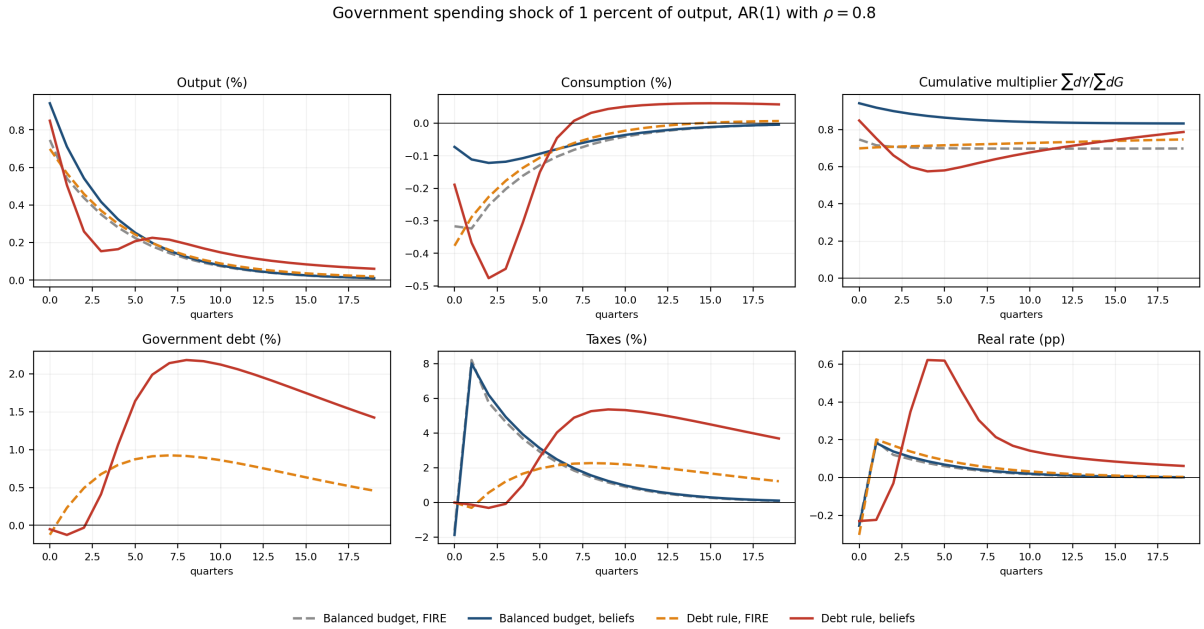


Figure 4

Government spending up one percent of output, AR(1) with $\rho = 0.8$. The multiplier panel plots the cumulative multiplier $\sum dY / \sum dG$; the period ratio is not informative once the shock has decayed.

The main result is that the belief friction raises the fiscal multiplier, having lowered the monetary one. Under a balanced budget the impact multiplier rises from 0.747 to 0.942 and the cumulative multiplier from 0.698 to 0.833; under the debt rule from 0.699 to 0.849 and from 0.747 to 0.787. This is the opposite sign from the effect of the same friction on monetary transmission in Table 3. The reason for this is that the two policies work through opposite margins. Monetary policy operates by inducing a private response, so a friction that delays that response weakens

	Impact	Cumulative	$C(0)$, %	Debt peak, %	$r(0)$, pp
Balanced budget, full information	0.747	0.698	−0.317		−0.244
Balanced budget, beliefs	**0.942**	**0.833**	−0.073		−0.255
Debt rule, full information	0.699	0.747	−0.377	0.92	−0.303
Debt rule, beliefs	**0.849**	**0.787**	−0.189	2.18	−0.229

Table 6

Output multipliers and impact responses to a government spending shock. $\vartheta = -0.19$, $\kappa_c = 0.23$, $\lambda = 0.75$, $\kappa = -0.21$.

it. Fiscal policy operates despite an offsetting private response, so the same delay postpones the crowding-out and strengthens it. The mechanism is visible in the consumption column: with beliefs, consumption falls 0.07 percent on impact instead of 0.32 percent.

Second, the results show that deficit financing raises the cumulative multiplier and lowers the impact multiplier. Under full information the impact multiplier falls from 0.747 to 0.699 when debt absorbs the shock, while the cumulative multiplier rises from 0.698 to 0.747. Postponing the tax burden helps at longer horizons, as one would expect. The impact result is less obvious and is specific to heterogeneous agents: the deficit raises the supply of government debt, and since the asset market must clear at $A_t = B_t$ with the real rate pinned by the policy rule, the additional debt has to be absorbed by higher household saving. Consumption therefore falls more on impact under the rule, 0.38 percent against 0.32, even though taxes have not yet risen. A representative agent model has no counterpart to this channel, because there the household's asset position is not a constraint on its consumption decision.

The largest debt response, 2.18 percent against 0.92, occurs when beliefs and the debt rule are combined, because the higher multiplier raises output and therefore the tax base only with a delay, so debt runs further before consolidation begins. This is the configuration in which the interaction between the financing rule and the belief distortion is strongest, and it is the one for which how passive is the rule matters the most: with $\phi_B \leq r$ the debt path in that configuration does not converge.

6 Conclusions

We estimate the belief process of a heterogeneous agent New Keynesian model from household survey data rather than calibrating it. The specification is general enough to nest several alternatives: households are pessimistic about their own employment risk, the distortion is ordered by the value of states, and a large anchoring component turns out to be survey response behaviour rather than a belief primitive.

I introduce this belief process in a macro model with uninsurable risk. At the estimated

values monetary transmission is dampened, peaks one quarter after impact rather than on impact, and is twenty percent stronger for contractions than for expansions. The results remain robust to the introduction of capital and related frictions. Furthermore, in a representative agent version the persistence induced by the belief process survive while the asymmetry vanishes, because the belief distortion is ordered by value differences that only uninsurable risk can make large. The two frictions are complements. The same friction raises the fiscal multiplier while lowering the monetary one, since monetary policy needs a private response that the friction delays while fiscal policy works despite one.

References

George-Marios Angeletos and Zhen Huo. Myopia and anchoring. *American Economic Review*, 111(4):1166–1200, 2021.

Adrien Auclert, Matthew Rognlie, and Ludwig Straub. The intertemporal keynesian cross. Working Paper 25020, National Bureau of Economic Research, 2018.

Adrien Auclert, Matthew Rognlie, and Ludwig Straub. Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model. Working Paper 26647, National Bureau of Economic Research, 2020.

Adrien Auclert, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408, 2021.

Bence Bardóczy and Joao Guerreiro. Unemployment insurance in macroeconomic stabilization with imperfect expectations. Manuscript, 2024.

Marzio Bassanin, Ester Faia, and Valeria Patella. Ambiguity attitudes and the leverage cycle. *Journal of International Economics*, 129:103436, 2021.

Johannes Boehm, Etienne Fize, and Xavier Jaravel. Five facts about MPCs: Evidence from a randomized experiment. *American Economic Review*, 115(1):1–42, 2025. doi: 10.1257/aer.20240138.

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Diagnostic expectations and credit cycles. *Journal of Finance*, 73(1):199–227, 2018.

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

Christopher D. Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, 2006.

Christopher D. Carroll, Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. Sticky expectations and consumption dynamics. *American Economic Journal: Macroeconomics*, 12(3):40–76, 2020.

Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45, 2005.

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

Roberto Colarieti, Pierfrancesco Mei, and Stefanie Stantcheva. The how and why of household reactions to income shocks. Working Paper 32191, National Bureau of Economic Research, 2024.

Martin S Eichenbaum, Joao Guerreiro, and Jana Obradovic. Ricardian non-equivalence. Technical report, National Bureau of Economic Research, 2026.

Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003.

Emmanuel Farhi and Iván Werning. Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, 109(11):3887–3928, 2019.

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

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

Lars Peter Hansen and Thomas J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001.

Per Krusell and Anthony A Smith, Jr. Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy*, 106(5):867–896, 1998.

Eric M. Leeper. Equilibria under active and passive monetary and fiscal policies. *Journal of Monetary Economics*, 27(1):129–147, 1991.

N. Gregory Mankiw and Ricardo Reis. Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics*, 117(4):1295–1328, 2002.

Filip Matějka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298, 2015.

Benjamin Moll. The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics. *The Economic Journal*, 136(676):1173–1216, 2026.

Jacob D. Orchard, Valerie A. Ramey, and Johannes F. Wieland. Micro MPCs and macro counterfactuals: The case of the 2008 rebates. *Quarterly Journal of Economics*, 140(3):2001–2052, 2025. doi: 10.1093/qje/qjaf015.

Oliver Pfäuti and Fabian Seyrich. A behavioral heterogeneous agent new keynesian model. Manuscript, 2026.

Ricardo Reis. Inattentive consumers. *Journal of Monetary Economics*, 53(8):1761–1800, 2006.

Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50 (3):665–690, 2003.

Thomas D. Tallarini. Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45 (3):507–532, 2000.

Silvana Tenreyro and Gregory Thwaites. Pushing on a string: US monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74, 2016.

A Recursivity of the household problem with distorted beliefs

The household's problem stated over infinite histories of idiosyncratic shocks is:

$$V(s_0, a_0) = \min_{\pi(s^t)} \max_{(c_t, n_t, a_t)} \sum_{t=0}^{\infty} \beta^t \sum_{s^t} \pi(s^t) u(c_t, n_t) + \frac{1-\beta}{|\vartheta|} \sum_{t=1}^{\infty} \beta^t \sum_{s^t} \pi(s^t) \ln \frac{\pi(s^t)}{\bar{\pi}(s^t)}, \tag{19}$$

subject to the budget and borrowing constraints and to $\sum_{s^t} \pi(s^t) = 1$, where $\bar{\pi}$ is the objective measure and the second term is the discounted Kullback-Leibler divergence of the subjective from the objective measure. This appendix shows that the sequence problem (19), in which the household chooses consumption, labour supply, assets and a subjective probability measure over the whole history of idiosyncratic shocks, is equivalent to a recursive problem in the two state variables (s, a_-) and that the optimal plan is dynamically consistent. It then verifies that

the two extensions used in the paper, the compression (11) and the dependence of the penalty parameter on the asset distribution, preserve these properties.

Notation follows the main text: s denotes the exogenous idiosyncratic state, $e(s)$ the labour productivity associated with it, a_- assets carried into the period, and $s^t = (s_1, \ldots, s_t)$ a history. The objective measure is $\bar{\pi}$ and the subjective measure π. We write $\bar{\mathbb{E}}$ for expectations under $\bar{\pi}$ and $\mathbb{E}^\pi$ for expectations under π. The relative entropy of π with respect to $\bar{\pi}$ on a finite set is:

$$\mathcal{K}\big(\pi \,\|\, \bar{\pi}\big) = \sum_s \pi(s) \ln \frac{\pi(s)}{\bar{\pi}(s)} \;\geq 0, \tag{20}$$

with equality if and only if $\pi = \bar{\pi}$. Let jus also clarify the relation between the penalty parameter used here and the tilt parameter used in the main draft. We can write the penalty coefficient as $\theta > 0$, so that the inner problem is:

$$\min_\pi \Big\{ \mathbb{E}^\pi[V] + \theta\, \mathcal{K}(\pi \,\|\, \bar{\pi}) \Big\}. \tag{21}$$

The tilt parameter used in the main text is:$\vartheta = -\frac{1}{\theta} < 0$, so that a small $|\vartheta|$ corresponds to a large penalty θ and hence to beliefs close to the objective measure. With this correspondence the penalty term in the sequence objective (19) carries the coefficient $(1-\beta)/|\vartheta| = -(1-\beta)/\vartheta$, which is positive. Note that if it were written instead with a coefficient $+(1-\beta)/\vartheta$, the penalty on deviations would be negative and the inner minimisation unbounded: the household could drive $\mathcal{K}$ to infinity and the objective to minus infinity.

The first step is a chain rule: the relative entropy of a measure over histories is the discounted sum of the relative entropies of its one-step conditionals.

Lemma 2 (Additive decomposition)**.** *Let π and $\bar{\pi}$ be measures over histories with $\pi \ll \bar{\pi}$, and factor them as $\pi(s^t) = \prod_{j=1}^t \pi(s_j \mid s^{j-1})$ and likewise for $\bar{\pi}$. Then:*

$$\frac{1-\beta}{|\vartheta|} \sum_{t=1}^\infty \beta^t \sum_{s^t} \pi(s^t) \ln \frac{\pi(s^t)}{\bar{\pi}(s^t)} \;=\; \frac{1}{|\vartheta|} \sum_{t=1}^\infty \beta^t\, \mathbb{E}^\pi \Big[\mathcal{K}\big(\pi(\cdot \mid s^{t-1}) \,\big\|\, \bar{\pi}(\cdot \mid s^{t-1})\big) \Big]. \tag{22}$$

Proof. Taking logs of the factorisation: $\ln \frac{\pi(s^t)}{\bar{\pi}(s^t)} = \sum_{j=1}^t \ell_j$ where $\ell_j \equiv \ln \frac{\pi(s_j|s^{j-1})}{\bar{\pi}(s_j|s^{j-1})}$. Hence:

$$\sum_{t=1}^\infty \beta^t\, \mathbb{E}^\pi \Big[\sum_{j=1}^t \ell_j \Big] = \sum_{j=1}^\infty \mathbb{E}^\pi[\ell_j] \sum_{t=j}^\infty \beta^t = \frac{1}{1-\beta} \sum_{j=1}^\infty \beta^j\, \mathbb{E}^\pi[\ell_j], \tag{23}$$

where the interchange of summation is justified by nonnegativity of each $\mathbb{E}^\pi[\ell_j]$. By the law of iterated expectations, $\mathbb{E}^\pi[\ell_j] = \mathbb{E}^\pi\big[\mathcal{K}(\pi(\cdot \mid s^{j-1}) \| \bar{\pi}(\cdot \mid s^{j-1}))\big]$. Multiplying by $(1-\beta)/|\vartheta|$ gives (22). □

The factor $(1-\beta)$ in (19) is exactly what cancels the $1/(1-\beta)$ generated by the double sum, leaving each date's conditional entropy discounted by β^t and no more.

Lemma 3 (The Inner Minimization Problem)**.** *Let $V : \mathcal{S} \to \mathbb{R}$ be bounded and $\theta > 0$. Then*

$$\min_{\pi \ll \bar{\pi}} \Big\{ \mathbb{E}^{\pi}[V] + \theta\, \mathcal{K}(\pi \| \bar{\pi}) \Big\} = -\theta \ln \bar{\mathbb{E}}\big[e^{-V/\theta}\big], \tag{24}$$

and the unique minimiser is:

$$\pi^*(s) = \frac{\bar{\pi}(s)\, e^{-V(s)/\theta}}{\bar{\mathbb{E}}\big[e^{-V/\theta}\big]}. \tag{25}$$

In terms of $\vartheta = -1/\theta < 0$, the minimised value is $\vartheta^{-1} \ln \bar{\mathbb{E}}[e^{\vartheta V}]$ and the minimiser is $\pi^* \propto \bar{\pi}\, e^{\vartheta V}$, which is (7).

Proof. The objective is strictly convex in π on the simplex, since $\mathcal{K}$ is strictly convex and $\mathbb{E}^{\pi}[V]$ is linear, so the minimiser is unique. Form the Lagrangian with multiplier λ on $\sum_s \pi(s) = 1$:

$$\mathcal{L} = \sum_s \pi(s) V(s) + \theta \sum_s \pi(s) \ln \frac{\pi(s)}{\bar{\pi}(s)} - \lambda \Big(\sum_s \pi(s) - 1 \Big). \tag{26}$$

The first-order condition with respect to $\pi(s)$ is:

$$V(s) + \theta \Big(\ln \frac{\pi(s)}{\bar{\pi}(s)} + 1 \Big) - \lambda = 0 \quad \Longrightarrow \quad \pi(s) = \bar{\pi}(s)\, e^{(\lambda - \theta)/\theta}\, e^{-V(s)/\theta}. \tag{27}$$

Imposing $\sum_s \pi(s) = 1$ pins the constant and gives (25). Substituting back:

$$\mathbb{E}^{\pi^*}[V] + \theta\, \mathcal{K}(\pi^* \| \bar{\pi}) = \sum_s \pi^*(s) \Big[V(s) + \theta \ln \frac{\pi^*(s)}{\bar{\pi}(s)} \Big] \tag{28}$$

$$= \sum_s \pi^*(s) \Big[V(s) + \theta \Big(- \frac{V(s)}{\theta} - \ln \bar{\mathbb{E}}\big[e^{-V/\theta}\big] \Big) \Big] = -\theta \ln \bar{\mathbb{E}}\big[e^{-V/\theta}\big], \tag{29}$$

where the terms in V cancel because the exponent in (25) is exactly $-V/\theta$ up to the normalising constant. Nonnegativity of the minimiser and absolute continuity are immediate from (25), since $\bar{\pi}$ has full support on $\mathcal{S}$. □

Lemma 3 is the source of the interpretation given in the main text: the minimiser is an exponential tilt of the reference measure, an Esscher transform, and the ratio $m(s) = \pi^*(s)/\bar{\pi}(s) = e^{\vartheta V(s)}/\bar{\mathbb{E}}[e^{\vartheta V}]$ satisfies $\bar{\mathbb{E}}[m] = 1$ by construction, so it is a Radon-Nikodym derivative and a martingale distortion.

Proposition 1 (Recursive representation)**.** *Let the budget correspondence be:*

$$\Gamma(s, a_-) = \Big\{ (c, n, a) \;:\; c + a = (1 + r) a_- + w\, e(s)\, n + T(s), \quad c > 0,\ n \geq 0,\ a \geq 0 \Big\}, \tag{30}$$

and suppose the state space $\mathcal{S} \times \mathcal{A}$ is compact and Γ is nonempty, compact-valued and continuous. Then the sequence problem (19) is solved by the fixed point of:

$$(\mathcal{T}W)(s, a_-) = \max_{(c,n,a) \in \Gamma(s, a_-)} \Big\{ u(c, n) + \frac{\beta}{\vartheta} \ln \bar{\mathbb{E}}\Big[e^{\vartheta W(s', a)} \,\Big|\, s \Big] \Big\}, \tag{31}$$

and the value of the sequence problem equals that fixed point.

Proof. By Lemma 2 the objective in (19) can be written:

$$\sum_{t=0}^{\infty} \beta^t \, \mathbb{E}^{\pi}\big[u(c_t, n_t)\big] + \frac{1}{|\vartheta|} \sum_{t=1}^{\infty} \beta^t \, \mathbb{E}^{\pi}\big[\mathcal{K}_t\big], \qquad \mathcal{K}_t \equiv \mathcal{K}\big(\pi(\cdot \mid s^{t-1}) \,\big\|\, \bar{\pi}(\cdot \mid s^{t-1})\big). \tag{32}$$

Both the felicity sum and the penalty sum are additively separable across dates with the common discount factor β, and the date-t penalty depends on the control $\pi(\cdot \mid s^{t-1})$ at that node only. The problem therefore has the standard additively separable dynamic programming structure, with the enlarged control $(c_t, n_t, a_t, \pi(\cdot \mid s^{t-1}))$ and the max and min taken jointly. Applying the principle of optimality yields:

$$V(s, a_-) = \max_{(c,n,a)\in\Gamma} \min_{\pi(\cdot|s)} \Big\{ u(c,n) + \beta \, \mathbb{E}^{\pi}\big[V(s', a)\big] + \frac{\beta}{|\vartheta|} \mathcal{K}\big(\pi(\cdot \mid s) \,\|\, \bar{\pi}(\cdot \mid s)\big) \Big\}. \tag{33}$$

The inner problem is exactly (21) with $\theta = |\vartheta|^{-1}$ scaled by β and with $V(\cdot, a)$ in the role of V; note that a is chosen before the minimisation and is therefore a constant in it. Lemma 3 evaluates it in closed form and gives (31) and (8). The order of max and min is immaterial: the objective is concave in (c, n, a) for concave u and convex in π, and both choice sets are convex and compact, so the minimax theorem applies and the saddle point exists. □

Proposition 2 (Dynamic consistency)**.** *The plan* $\{c_t, n_t, a_t, \pi^*(\cdot \mid s^{t-1})\}$ *that solves* (19) *at date* 0 *remains optimal for the problem restarted at any date* t *and any history* s^t*. Equivalently, the household has no incentive to revise its subjective measure, and no commitment technology is required.*

Proof. By Lemma 2 the penalty is a sum of one-step conditional entropies, each attached to a single node of the event tree and each entering with the same coefficient $1/|\vartheta|$, independent of t and of the history. The date-t continuation of the date-0 objective is therefore, node by node, the date-t objective. By Lemma 3 the minimiser at a node depends on the continuation value at the successor nodes and on nothing else, and by Proposition 1 that continuation value depends on the history only through (s_t, a_{t-1}). Hence the date-0 minimiser restricted to the subtree rooted at s^t coincides with the date-t minimiser, and the same argument applied to the maximisation gives the result. □

Note that dynamic consistency is a property of the multiplier formulation, not of robust preferences in general. If instead the household were given a fixed budget of relative entropy to allocate across dates, as in a constraint formulation, the date-0 allocation of that budget would generally not be optimal from the perspective of date t, and the problem would fail to be recursive unless the constraint set is rectangular. Fixing a constant ϑ makes the implied constraint set rectangular by construction, and this is the reason the paper works with the multiplier form.

At last, we check if recursivity survives when we extent the belief process to include the compression and the heterogeneity with respect to assets.

Corollary 1 (Compression)**.** *Replace the reference measure* $\bar{\pi}(\cdot \mid s)$ *by the mixture* $\hat{\pi}(\cdot \mid s) = (1-\kappa_c)\bar{\pi}(\cdot \mid s) + \kappa_c\, p(\cdot)$ *of* (11)*, where* p *is the stationary distribution and* $\kappa_c \in [0,1)$*. Then Lemmas 2 and 3 hold verbatim with* $\bar{\pi}$ *replaced by* $\hat{\pi}$*.*

Proof. $\hat{\pi}(\cdot \mid s)$ is a probability measure on $\mathcal{S}$ for every s, and it is Markov in s because $\bar{\pi}$ is and p does not depend on s. It has full support whenever $\bar{\pi}$ or p does, so $\pi \ll \hat{\pi}$ is satisfied by any candidate. None of the proofs uses any property of the reference measure beyond these. Note that the distribution of the economy continues to evolve under the objective $\bar{\pi}$: $\hat{\pi}$ enters the household's problem only, which is what makes the household wrong about its environment rather than changing the environment. □

Corollary 2 (Asset-dependent penalty)**.** *Let the penalty depend on assets,* $\vartheta(a) = \vartheta + \gamma\big(x(a) - \frac{1}{2}\big)$ *with* $x(a) = a/(a + a_{\text{ref}})$*, and suppose* $\vartheta(a) < 0$ *for all* $a \in \mathcal{A}$*. Then Proposition 1 holds with* ϑ *replaced by* $\vartheta(a)$ *in* (31)*.*

Proof. In (33) the asset choice a is made before the inner minimisation, so $\vartheta(a)$ is a constant in that problem and Lemma 3 applies pointwise in a with $\theta = |\vartheta(a)|^{-1}$. □

B Numerical transformation to include the distorted beliefs

The model is solved with the sequence-space Jacobian method of Auclert et al. [2021] and the endogenous gridpoint method of Carroll [2006]. Implementing the belief distortion inside that framework requires a transformation of the household problem, and because the reason is not obvious we set it out in full.

The envelope conditions in the SSJ toolkit. The primitive object of a heterogeneous agent block is a function that performs one step of backward induction: it receives the date $t+1$ marginal value function and returns the date t policies and marginal value function. Arguments carrying the suffix p denote next-period objects, so the function receives V_a^p and returns V_a. The toolkit iterates this function to convergence to obtain the stationary policy rules, and differentiates it along a perturbed price path to obtain the Jacobians.

The toolkit applies, across iterations the expectation itself, so what the function actually receives is:

$$V_a^p(s,a) \;=\; \sum_{s'} \bar{\Pi}(s,s')\, V_a(s',a), \tag{34}$$

the objective expectation, already summed over next-period states.

We apply a transformation that allows us to introduce the tilt. Notice that (**??**) is a ratio of two objective expectations. Both the numerator and the denominator are of the form $\sum_{s'} \bar{\Pi}(s, s')\, g(s', a)$, which is precisely what the toolkit computes. It is therefore enough to arrange for the single available application of $\bar{\Pi}$ to land on the exponentiated objects rather than on V_a. We do this by carrying three backward variables:

$$X = e^{\vartheta(V - V_0)}, \qquad Y = e^{\vartheta(V - V_0)}\, V_a, \qquad \mathcal{V} = V, \tag{35}$$

so that the toolkit returns:

$$X^p(s, a) = \sum_{s'} \bar{\Pi}(s, s')\, e^{\vartheta(V(s', a) - V_0)}, \qquad Y^p(s, a) = \sum_{s'} \bar{\Pi}(s, s')\, e^{\vartheta(V(s', a) - V_0)} V_a(s', a), \tag{36}$$

and the two objects needed in (**??**) are available directly:

$$\tilde{\mathbb{E}}\big[V_a'\big] = \frac{Y^p}{X^p}, \qquad W = \frac{1}{\vartheta} \ln X^p + V_0, \tag{37}$$

where W is the risk-sensitive certainty equivalent of (8). The third variable $\mathcal{V}$ carries V itself forward so that the exponentials in (35) can be rebuilt at the next step. The constant V_0 cancels from the ratio and from (37), and serves only to recentre the exponent: since V is a level that scales with $1/(1 - \beta)$, working with $V - V_0$ keeps $e^{\vartheta(V - V_0)}$ away from overflow. In the implementation X is stored as expm1 $\big(\vartheta(V - V_0)\big)$ and decoded with log1p, which removes the cancellation error that would otherwise appear as $\vartheta \to 0$ and makes the rational expectations limit exact rather than merely close.