

# Perceived Income Risk Across the Income Distribution

## Abstract

Low-income workers report wage uncertainty about twice as large as high-income workers, conditional on staying in the same job, yet matched administrative spells show no comparable gradient in realised risk, so the gap is one of beliefs. I show that a single constant-gain learning rule organises three features of perceived risk by income: its level, its persistence, and its response to realised shocks. Estimated within person, the bottom updates most slowly, reads the largest share of each shock as permanent, and carries a prior far above its realised transitory risk. I microfound the gradient in how pay is delivered. Perceived risk departs from realised risk in opposite directions across the distribution, above the truth at the bottom and below it at the top, a pattern no one-directional friction can produce. Two observable features of pay, the noisy reading of fragmented low-income pay and the institutional labelling of high-income transitory pay, reproduce all three estimated patterns with no behavioural bias. Embedded in an Aiyagari economy, the calibrated beliefs open precautionary-saving and marginal-propensity-to-consume gradients that a full-information calibration cannot generate; the wedge falls on the unconstrained middle and top, not on the bottom that holds the most distorted beliefs.

## 1 Introduction

In the New York Fed’s Survey of Consumer Expectations (SCE), workers in households earning under \$50k report wage uncertainty about 1.8 times as large as workers in households above \$100k, conditional on staying in the same job. Their beliefs are also more persistent, and they revise them more strongly after realised wage shocks. Matched same-job spells in the Survey of Income and Program Participation (SIPP) show no comparable gradient in the realised wage process; if anything realised risk rises mildly with income. The gradient is therefore one of beliefs, not of realised risk, and the standard heterogeneous-agent practice of setting perceived risk equal to realised risk fails on all three margins at once.

The paper proceeds in three steps. The first is descriptive: perceived wage risk falls steeply with income, is most persistent at the bottom, and co-moves with realised shocks (Sections 3–4). The second fits a learning rule, in which agents learn the variance of their permanent earning power from realised pay signals through an exponentially weighted average of squared surprises, the steady-state Kalman filter for an unknown variance. The rule has three interpretable parameters, how fast beliefs update, what share of each squared surprise is read as permanent, and a prior on transitory risk, and estimated within person and by income group the bottom updates most slowly, reads roughly three times as much of each shock as permanent, and holds a prior almost four times its realised transitory risk, while the middle and top sit close to one another and near their realised processes (Sections 5–7). The persistence and prior gradients are significant; the attribution gradient is monotone but not.

The third step asks why the deviations take the form they do, which is where the paper departs from the standard model. A belief rule with an update speed and an attribution share already reproduces the persistence and the level of the gradient, but it does not explain them, since those are reduced-form objects estimated freely per group. What they leave unexplained runs in both directions at once: measured against the risk each group actually faces, the bottom *over*-perceives transitory risk while the middle and top *under*-perceive it, in the same cross-section. No single friction can produce a gradient that crosses its own benchmark this way, because the two candidate frictions are each one-sided. Noise in reading one’s own pay can only raise perceived risk, being the observational twin of a transitory shock, while recognising part of one’s pay as transitory and netting it out can only lower it. The data therefore require one force of each sign, and Appendix F shows that neither can be dropped without rejection.

Both forces are observable features of how pay is delivered. Pay at the bottom arrives in noisy fragments, variable shift counts, cash tips, and irregular supplements, which makes a worker’s own reading of her wage change noisy and inflates her perceived risk. Pay at the top arrives with its transitory component labelled, as bonuses and commissions on separate line items, which the worker nets out and which deflates her perceived risk. A Bayesian agent facing this signal structure reproduces all three estimated patterns with no behavioural bias: the prior and persistence directly, and the attribution once the cell-mean attenuation of the regressor is taken into account (Section 8). The income gradient in beliefs is a gradient in how pay is delivered, not in how risky pay is or in how agents react to it.

Embedded in an Aiyagari buffer-stock economy with group-specific liquidity, the calibrated beliefs open precautionary-saving and marginal-propensity-to-consume (MPC) gradients that a full-information calibration cannot generate (Sections 9–10). The wedge falls on the middle and top: they under-perceive the permanent risk that drives saving, are unconstrained, and so hold less wealth and consume more of a windfall. The bottom holds the most distorted beliefs, but its distortion sits in the transitory prior rather than in the permanent belief that drives saving, and being near hand-to-mouth it is moved by liquidity rather than by its own beliefs.

*Relation to the literature.* Wang (2023) documents the level gap between perceived and calibrated income risk in the SCE; I add its cross-sectional structure, its within-person dynamics, an estimated updating rule, and a mechanism. Rozsypal and Schlafmann (2023) document income-graded bias in expected income *levels*; I study the second moment. Caplin et al. (2023) find perceived earnings risk below econometric estimates in Danish register-linked data, consistent with the under-perception I find higher up the distribution. Kaufmann and Pistaferri (2009) show that superior information rationalises part of such gaps; the labelled-pay channel here is that information channel made observable and graded by income. Broer et al. (2022) study information heterogeneity in heterogeneous-agent settings. The mechanism rests on documented compensation facts: performance pay concentrates at the top (Lemieux et al., 2009), high-income shocks are predominantly transitory (Guvenen et al., 2021), bottom-of-distribution income is dominated by volatile, fragmented hours and gig flows (Hardy and Ziliak, 2014; Morduch and Schneider, 2017), and recessions hit the bottom hardest (Hoynes et al., 2012).

The empirical results (Sections 3–7, Appendices A–F) are complete. The mechanism of Section 8 is calibrated and directly testable. The model of Section 9 is solved in partial equilibrium; general

equilibrium is left for a two-asset extension.

## 2 Data and measurement

*Samples.* The SIPP sample covers reference years 2013–2023 and the SCE runs 2013–2025; both are monthly within-respondent panels. I restrict to same-job, same-employer spells of at least 6 consecutive months and apply survey weights throughout. Respondents are sorted into three income groups by median annual household income: Bot ( $< \$50k$ ), Mid ( $\$50k$ – $\$100k$ ), Top ( $\geq \$100k$ ).<sup>1</sup>

*SCE.* The density question asks for the distribution of next-12-month earnings growth “*on the same job, same hours, same employer*”, on a 10-bin histogram. The object I take from it, perceived variance  $\bar{\sigma}_{\text{perc}}^2$ , is the bin-midpoint variance of each respondent’s reported histogram. Respondents placing all mass in a single bin are excluded, and the result is trimmed at 1%/99%. I residualise on age, sex, state, education, survey tenure, and year-month fixed effects; survey tenure absorbs panel conditioning (Kim and Binder, 2023). I drop each respondent’s observations from the first wave in which they report a job change.

*SIPP.* Per-month earnings include regular wage and salary plus bonuses, commissions, tips, and overtime; the self-employed are excluded. I residualise log wages on the same demographics, state, and year-month fixed effects and trim squared changes at 1%/99%.


## 3 Perceived risk by income group

### Levels

SCE perceived variance is steeply income-graded throughout the sample (Figure 1). The weighted group means are 20.6 at the bottom, 11.3 in the middle, and 11.4 at the top, so perceived variance at the bottom is about 1.8 times that at the top, and equality of bottom and top is rejected. The gap is a step at the bottom: the middle and top are statistically indistinguishable and track each other closely, while the bottom sits about 9 points above both. The bottom series also drifts upward over time, possibly reflecting the fixed nominal cutoffs or greater cyclical exposure.

---

<sup>1</sup>The cutoffs are nominal and fixed across the sample. The SCE skews toward higher incomes, so group sizes are uneven (estimation sample: 1,531 / 7,501 / 11,331 observations).



**Figure 1:** Group-mean perceived variance  $\bar{\sigma}_{\text{perc}}^2$  in  $(\%)^2$ , quarterly with a 4-quarter trailing MA. Matched-stayer SCE panel, residualised as in Section 2.

### Persistence

The level gradient is silent about dynamics. I estimate a within-respondent AR(1) for  $\bar{\sigma}_{\text{perc}}^2$ , with persistence coefficient  $\hat{\rho}_g$ , by Blundell-Bond (BB) system GMM (Blundell and Bond, 1998). The short, wide SCE panel makes the within estimator severely Nickell-biased, and BB instruments the lagged level with deeper lags to remove that bias. I use BB throughout the paper, following Ma et al. (2020), who work in the same short-panel setting.

Group	Monthly		Quarterly	
	$\hat{\rho}_g$	SE	$\hat{\rho}_g$	SE
Bot	+0.343***	0.046	+0.632***	0.062
Mid	+0.142***	0.026	+0.594***	0.088
Top	+0.144***	0.025	+0.377***	0.076
Wald ( $B = T$ )	$p = 0.0001$		$p = 0.009$	

**Table 1:** Within-respondent AR(1) persistence  $\hat{\rho}_g$  of perceived variance, Blundell-Bond system GMM. Instruments:  $\bar{\sigma}_{\text{perc}}^2$  at lags 3–4 in differences plus the BB level moment, collapsed. Two-way clustered SEs (respondent, quarter). The same instrument set is used at both cadences.

Perceived variance at the bottom is the most persistent of the three groups, and equality with the top is rejected at both cadences. The quarterly estimates exceed the monthly ones for every group, which time aggregation of an AR(1) cannot produce; the likely cause is transitory response noise in the monthly reports, which attenuates the monthly coefficients.

One main takeaway from this exercise is that the standard heterogeneous-agent assumption, in which perceived risk is set equal to realised risk, does not hold. That practice calibrates  $\sigma^2$  from realised wage data, but it matches neither the level nor the persistence of perceived risk documented here, and the next section shows the realised process cannot supply the gradient either.

## 4 The realised wage process

If the belief gradient reflected a gradient in actual risk, it would be visible in the realised wage process. It is not. Realised shocks nonetheless move beliefs, and together these two facts pin down what a belief rule must deliver.

### Wage model

I use the standard permanent-transitory decomposition (Meghir and Pistaferri, 2004):

$$w_{i,t} = z_{i,t} + e_{i,t}, \quad e_{i,t} = p_{i,t} + \theta_{i,t}, \quad p_{i,t} = p_{i,t-1} + \psi_{i,t}.$$

Here  $w_{i,t}$  is the residualised log wage,  $z_{i,t}$  the predictable Mincerian path, and  $e_{i,t}$  the stochastic component that is the focus of the rest of the paper. The permanent shock  $\psi_{i,t}$  is iid with variance  $\sigma_\psi^2$  and accumulates through the random walk  $p_{i,t}$  (a merit raise or cost-of-living adjustment); the transitory shock  $\theta_{i,t}$  is iid with variance  $\sigma_\theta^2$  (a discretionary bonus or profit share).

### Matching SCE to the model

The SCE question conditions on *same job, hours, employer*, which fixes  $\Delta z = 0$ , so the elicited variance is that of the stochastic component alone:

$$\bar{\sigma}_{\text{perc},i,t}^2 \equiv \text{Var}_{i,t}(\Delta w_{i,t+1} \mid \Delta z = 0) = \text{Var}_{i,t}(\Delta e_{i,t+1}).$$

On SIPP I residualise (removing  $z$ ) and difference to recover  $\Delta e$  directly. The recorded wage is the amount paid per usual hour, so it includes tips, overtime, and any gap between actual and usual hours: the recorded counterpart of the worker's paycheck. Both sides therefore measure  $\text{Var}(\Delta e)$ .

### Identification of the wage primitives

The one-period change is  $\Delta e_t = \psi_t + \theta_t - \theta_{t-1}$ . Write  $\bar{S}_m$  for the variance of this monthly change,  $\mathbb{E}[(\Delta e)^2]$ . Because  $\psi$  and  $\theta$  are independent and iid, it has only two non-trivial moments per group:

$$\bar{S}_m = \sigma_\psi^2 + 2\sigma_\theta^2, \quad \text{Cov}(\Delta e_t, \Delta e_{t-1}) = -\sigma_\theta^2.$$

The lag-1 autocovariance identifies  $\sigma_\theta^2$  and the remaining variance identifies  $\sigma_\psi^2$ , so the two wage primitives are exactly identified per group.

Group	$\sigma_\theta^2$ (transitory)	$\sigma_\psi^2$ (permanent)	$\bar{S}_m$
Bot	2.52	11.92	16.96
Mid	4.42	11.61	20.45
Top	5.83	11.23	22.89

**Table 2:** SIPP permanent-transitory identification per income group, monthly cadence, in  $(\%)^2$ . Broad stayer panel, January seam dropped.

Permanent variance is roughly flat across groups ( $\approx 11.5$ ), while transitory variance rises mildly

with income, plausibly a bonus channel at the top. The realised gradient therefore runs, if anything, *opposite* to the perceived one.

### The FIRE benchmark

To compare the monthly process with the 12-month SCE horizon, I roll it forward twelve months under iid shocks (derivation in Appendix B):

$$\text{Var}(e_{t+12} - e_t) = 12\sigma_\psi^2 + 2\sigma_\theta^2, \quad (1)$$

which is what a full-information rational-expectations (FIRE) agent who knows the parameters would report.

Group	FIRE benchmark ( $12\sigma_\psi^2 + 2\sigma_\theta^2$ )	SCE observed $\bar{\sigma}_{\text{perc}}^2$	ratio
Bot	148.1	20.6	7.2×
Mid	148.2	11.3	13.1×
Top	146.4	11.4	12.8×

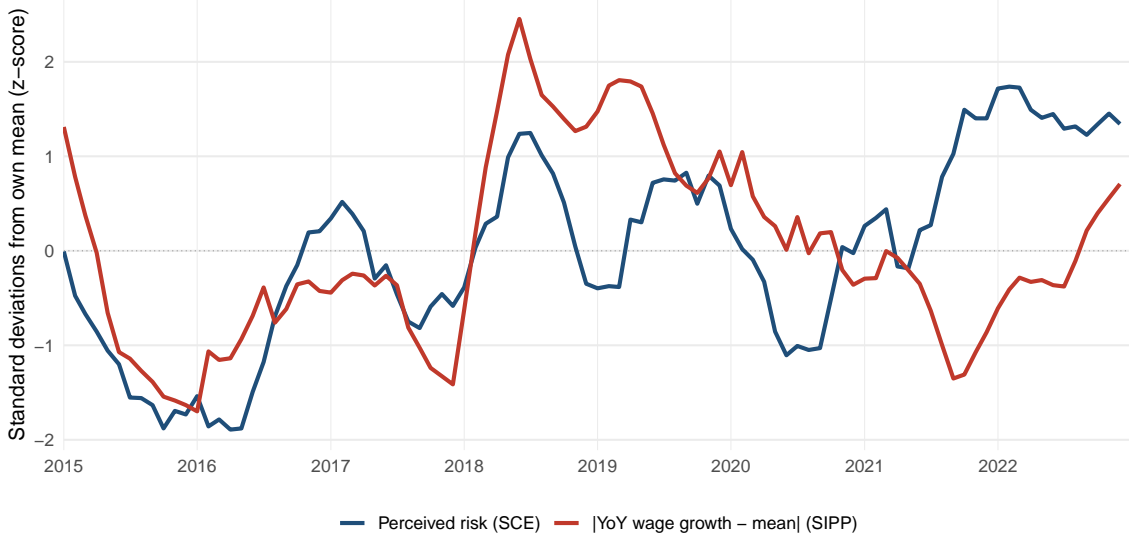
**Table 3:** FIRE-implied 12-month variance versus SCE-reported perceived variance, in (%)<sup>2</sup>.

*Reading the level gap.* These ratios overstate miscalibration, for a reason mechanical in the benchmark. Equation (1) scales permanent variance by twelve but transitory variance only by two, so any contamination loading onto  $\sigma_\psi^2$  is magnified twelvefold. SIPP earnings are self-reported, and reporting error is exactly such a contaminant: its classical component differences into an MA(1) and is read as transitory, while any persistent component is indistinguishable from a true permanent shock and is then multiplied by twelve (Bound and Krueger, 1991). Administrative payroll data tell the same story, placing job-stayers’ base-wage dispersion well below the benchmark (Grigsby et al., 2021). I therefore place no weight on the absolute ratios. The cross-group contrast is robust: error common across groups cancels in the comparison and cannot, by itself, give the two gradients opposite signs.

### Realised shocks still move beliefs

The cross-section rules out the realised process as the *source* of the gradient, but the realised process can still drive belief *dynamics*. Pooled perceived risk co-moves with the absolute deviation of aggregate wage growth from its sample mean (Figure 2), consistent with surprises to wage growth moving perceived risk over time.<sup>2</sup>

<sup>2</sup>Absolute deviations, since any realised shock should raise perceived uncertainty; squared deviations give a very similar plot. Section 7 splits shocks by sign.



**Figure 2:** SCE perceived risk and the absolute deviation of SIPP aggregate wage growth from its sample mean, both standardised, pooled across income groups, 2015–2022.

Together these facts impose two requirements on any belief rule. It must let perceived variance move, since the dynamics rule out constant beliefs; and it must let realised shocks drive that movement, with the response to the largely common signal differing by group.

## 5 A belief-updating rule

The agent cannot observe the permanent and transitory shocks separately. Each period she observes one squared pay change, which mixes the two ( $\mathbb{E}[(\Delta e)^2] = \sigma_\psi^2 + 2\sigma_\theta^2$ ), and must attribute the surprise, assigning a share  $\lambda$  to permanent earning power and the remainder to transitory pay. The SCE delivers one perceived-variance series per group, enough to identify one dynamic belief, not two. I let the permanent belief  $\hat{\sigma}_{\psi,t}^2$  evolve and treat the transitory belief as a static prior  $\sigma_{\theta,\text{perc}}^2$ . The permanent component is the welfare-relevant one, so it is the natural place for the dynamics; the cost is that any belief-side misattribution is absorbed into  $\lambda$ .

The permanent-variance belief follows an exponentially weighted moving average (EWMA) of past squared shocks, with the attribution share as a separate weight:

$$\hat{\sigma}_{\psi,t}^2 = \underbrace{(1 - \mu)}_{\text{persistence}} \hat{\sigma}_{\psi,t-1}^2 + \underbrace{\mu}_{\text{update weight}} \underbrace{\lambda}_{\text{attribution share}} (\Delta e_{t-1})^2, \quad (2)$$

where  $\mu$  is the update speed (a higher  $\mu$  moves the belief more on each period’s news) and  $\lambda$  is the share of each squared surprise read as informative about permanent earning power. Because the weights  $(1 - \mu)$  and  $\mu$  sum to one, (2) is the integrated GARCH(1,1) recursion in the squared surprise, which is also the steady-state Kalman filter for an unknown permanent variance modelled as a random walk:  $\mu$  plays the role of the steady-state Kalman gain. Appendix F gives the filter and shows how the gain is set by the signal-to-noise ratio. The agent is Bayesian given this information structure and applies the steady-state gain each period; she neither over- nor under-reacts relative to that benchmark (Section 8 returns to this point).

Substituting the belief into the FIRE identity (1) gives the level the agent reports:

$$\bar{\sigma}_{\text{perc},t}^2 = 12 \hat{\sigma}_{\psi,t}^2 + 2\sigma_{\theta,\text{perc}}^2. \quad (3)$$

Strict FIRE is the corner  $\mu = 0$  with both beliefs at the truth, giving a constant reported variance that the dynamics of Section 3 rule out. Within the learning family ( $\mu > 0$ ), FIRE survives as a long-run benchmark: taking unconditional expectations of (2) in a stationary state gives  $\mathbb{E}[\hat{\sigma}_{\psi}^2] = \lambda \bar{S}_m$ , so the belief is correct on average exactly when  $\lambda^* = \sigma_{\psi}^2 / \bar{S}_m$ , which Table 2 implies equals 0.70, 0.57, and 0.49 at the bottom, middle, and top, with the transitory condition  $\sigma_{\theta,\text{perc}}^2 = \sigma_{\theta}^2$  alongside. Section 7 shows neither holds.

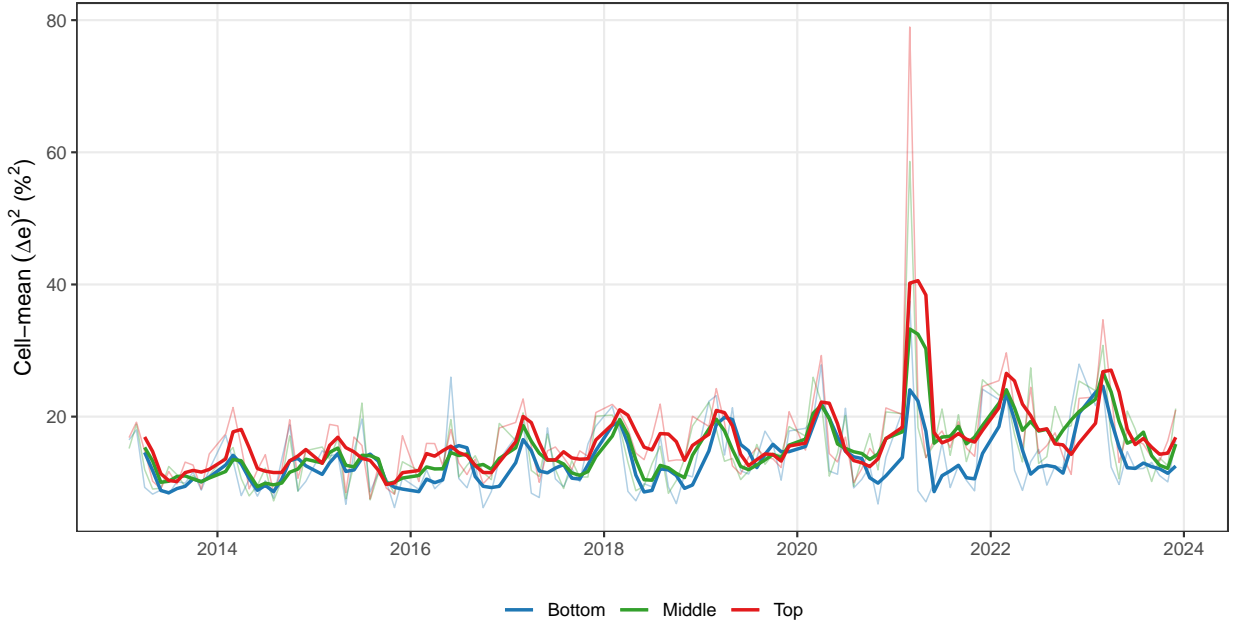
## 6 Estimation

The belief recursion (2) is written in the unobserved  $\hat{\sigma}_{\psi}^2$ . Combining it with the level identity (3) eliminates the belief and leaves an AR(1) in the *observable* perceived variance (Appendix A):

$$\bar{\sigma}_{\text{perc},i,t}^2 = \underbrace{2\mu_g \sigma_{\theta,\text{perc},g}^2}_{\beta_{\text{cons}}} + \underbrace{(1 - \mu_g)}_{\beta_{\text{lag}}} \bar{\sigma}_{\text{perc},i,t-1}^2 + \underbrace{12\mu_g \lambda_g}_{\beta_{\text{shock}}} (\Delta e_{i,t-1})^2 + \alpha_i + u_{i,t}, \quad (4)$$

with  $\alpha_i$  a person fixed effect. The three coefficients map one-to-one to the structural parameters:  $\mu_g = 1 - \hat{\beta}_{\text{lag}}$ ,  $\sigma_{\theta,\text{perc},g}^2 = \hat{\beta}_{\text{cons}} / (2\mu_g)$ , and  $\lambda_g = \hat{\beta}_{\text{shock}} / (12\mu_g)$ . Each is identified from a distinct feature of the data:  $\beta_{\text{lag}}$  from the within-respondent persistence,  $\beta_{\text{cons}}$  from the group level once the news loading is netted out, and  $\beta_{\text{shock}}$  from the response to the lagged realised shock.

*The shock regressor.* The rule calls for the agent's own  $(\Delta e_{i,t-1})^2$ . The SCE contains no individual wage histories, so I substitute the SIPP cell mean over workers in the same cell  $c = (\text{group} \times \text{Census division} \times \text{month})$ , following the JMP of Guerreiro (2023); workers in the same local labour market share components of wage news, transmitted through coworkers and local media and salient for expectations (Kuchler and Zafar, 2019). The cell series is smoothed with a 3-month trailing MA.



**Figure 3:** Cell-mean  $(\Delta e_{i,t})^2$  by income group, national monthly, in  $(\%)^2$ . Thin lines: raw monthly cell means. Thick lines: the 3-month trailing MA used as the regressor in (4).

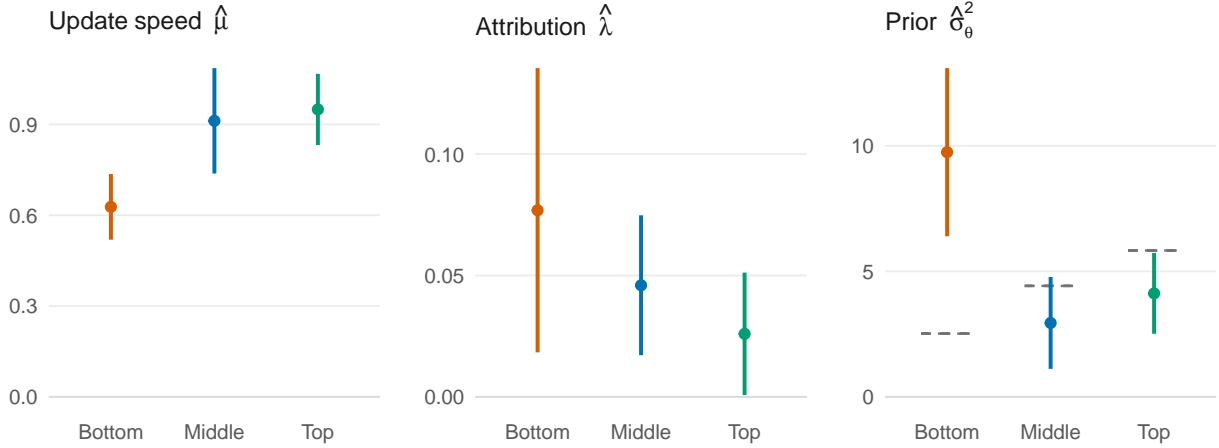
The top has the largest shock amplitude, consistent with its higher  $\sigma_\theta^2$  and with discretionary pay concentrating at the top (Güvener et al., 2021; Halvorsen et al., 2024). Because the bottom’s shocks are no larger than the top’s, any stronger belief response at the bottom is a property of beliefs, not of shocks.

*Attenuation of the shock loading.* Replacing the own shock with a cell mean is a measurement-error problem that attenuates  $\beta_{\text{shock}}$ , and hence  $\hat{\lambda}$ , toward zero. The attenuation works through finite-cell sampling noise: the cell mean tracks only the common component of a worker’s shock, so its informativeness about her own shock depends on how large the common component is and on the cell size. Appendix F measures the common-shock share in SIPP at about two percent of the variance of the squared change, which with the cell sizes used here implies an attenuation of roughly five to one, close to common across groups. The level of  $\hat{\lambda}$  is therefore not interpretable on its own; the calibration of Section 9 recovers the structural level by reproducing this attenuation inside a simulated cell merge.

*Instruments.* I instrument with  $\bar{\sigma}_{\text{perc}}^2$  at lags 3–4 and the shock at lags 1–3, together with BB level moments. The shock is treated as endogenous: the contemporaneous shock and the belief share a survey-timing and aggregate-news channel, and the Hansen  $J$  test rejects the contemporaneous moment set but accepts lags 1–3.

## 7 Belief estimates by income group

Joint GMM on (4) by income group yields the estimates in Figure 4, set against the realised SIPP transitory variance. The full estimates, with standard errors, the Hansen  $J$ , and sample sizes, are in Table 6 (Appendix D).



**Figure 4:** Headline estimates by income group: update speed  $\mu_g$ , attribution  $\lambda_g$ , and transitory prior  $\sigma_{\theta, \text{perc}, g}^2$ , with 95% confidence intervals. Grey dashes in the prior panel mark realised transitory variance  $\sigma_\theta^2$  (SIPP). The attribution is identified in cross-group ratios, so its panel shows slope, not level.

*Update speed.* The bottom updates significantly more slowly than the middle or top, matching the persistence ordering of Table 1. The influence of a given wage shock on its belief therefore decays more slowly.

*Attribution.* Each  $\hat{\lambda}_g$  is positive and significant, and the point estimates fall monotonically with income, the bottom's being three times the top's, although cross-group equality is not rejected. A common wage signal generates larger revisions of the permanent belief at the bottom, in point estimates if not yet in significance.

*Transitory prior.* Realised transitory variance rises with income; the perceived prior does the opposite. The bottom's prior sits almost four times its realised value, while the middle and top lie near or below theirs. The bottom is the only group whose prior is materially detached from its realised data, and the gradient is significant. The inflated prior is the dominant driver of the *level* of the bottom's perceived variance, separate from the news channel.

*Robustness.* Adding an occupation  $\times$  month fixed effect to the shock construction preserves the persistence and attribution gradients and sharpens the latter, while the prior gradient weakens (Appendix D.1); I do not adopt it as the baseline, since low-income workers sort into structurally more uncertain occupations (Morduch and Schneider, 2017), so partialling out occupation removes part of the channel. Appendix D.2 varies the instrument lags, the trim, and the COVID handling one at a time: the persistence gradient survives every variant except a tighter within-group trim, and the attribution gradient stays monotone throughout. When the shock is split by sign and groups are pooled for power (Appendix D.3), both components load positively, but upside news carries about twice the weight of downside (0.055 versus 0.027). My interpretation is that workers read good news as informative about permanent earning power and dismiss bad news as transitory, although this is the opposite of what loss aversion would predict.

## 8 A pay-delivery mechanism

The estimates leave the three parameters per group structurally unconnected. To connect them, I conjecture that two findings from the literature on how pay is delivered can explain their shape: low-income pay arrives in fragmented, hard-to-read form, while transitory pay at the top arrives institutionally labelled. Each maps to one observable primitive, and I show that both are necessary.

### Why two primitives, and which

The shape of the prior is what the mechanism has to reproduce. The perceived transitory prior lies above realised transitory variance at the bottom (9.75 against 2.52) and below it at the middle and top (2.94 against 4.42, 4.12 against 5.83), so perceived risk crosses its own realised benchmark. Two features of how pay is delivered can move the prior, and each moves it in only one direction: noise in reading one’s own pay can only raise it, and recognising part of one’s pay as transitory and netting it out can only lower it.

*Noisy reading of own pay.* The worker does not read her wage process from a clean statement. Pay at the bottom arrives as a stream of fragments: variable shift counts, cash tips, and irregular supplements, on top of hours that are themselves uncertain. With a variable schedule and irregular overtime, she cannot cleanly separate a change in her wage rate from a change in hours worked, so even her reading of what she earned per hour is noisy. I show its empirical counterpart in Appendix E, where the within-person hours volatility falls with income, which is consistent with reading noise largest at the bottom. I model her real-time reading as

$$\tilde{y}_{i,t} = e_{i,t} + \nu_{i,t}, \quad \nu_{i,t} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \omega_g^2), \quad (5)$$

a signal-extraction structure in the style of Pischke (1995), where  $\nu_{i,t}$  is *perception* noise of variance  $\omega_g^2$ , not income variation.<sup>3</sup> She feeds  $(\Delta \tilde{y}_t)^2$  into (2) as if it were  $(\Delta e_t)^2$ . A differenced reading error has the same variance-and-autocovariance signature as a true transitory shock, so it is recorded as additional transitory risk and raises the prior.

*Labelled transitory pay.* High-income compensation includes visibly transitory components: bonuses and commissions arrive as separate line items (Lemieux et al., 2009), and high-income shocks are predominantly transitory (Guiso et al., 2021). Low-income compensation is hours times wage plus tips, with no flag separating the transitory portion. The same logic applies to macro exposure: the top’s tends to arrive classifiably (furloughs with return dates, cyclical hours cuts), the bottom’s ambiguously (sector closures, the disappearance of gig work) (Hoynes et al., 2012). Let  $\rho_g \in [0, 1]$  be the fraction of transitory pay that arrives institutionally labelled. The labelled share is netted out before the update, so only the unlabelled residual feeds (2) and the prior falls.

Because each force is one-sided, a prior that crosses its benchmark cannot come from either alone. Appendix F confirms it: imposing either primitive to zero is rejected, each failing on the side of the crossing it cannot reach. The data require one force of each sign, and the income gradient is

---

<sup>3</sup>Recorded SIPP earnings capture the amounts actually paid, so  $\nu$  does not enter the moments of Section 4. The claim is that fragmented delivery makes the worker’s own reading of her pay noisy.

which force dominates where.

### From primitives to estimated gradients

The two primitives map onto the three estimated parameters:  $\rho_g$  governs the update speed and the attribution slope,  $\omega_g^2$  the transitory prior.

*Transitory prior.* The two forces pull in opposite directions, as above. Reading noise pushes the prior up: differencing the noisy reading (5) leaves  $\nu_t - \nu_{t-1}$ , an MA(1) with the variance-and-autocovariance signature of a true transitory shock, which the agent records as risk.<sup>4</sup> Labelling pushes the prior down: the flagged share is netted out, leaving  $(1 - \rho_g)\sigma_\theta^2$ . The prior is the sum,

$$\sigma_{\theta,\text{perc},g}^2 = (1 - \rho_g)\sigma_\theta^2 + \omega_g^2. \quad (6)$$

Noise dominates at the bottom (high  $\omega^2$ , low  $\rho$ ) and the prior sits above realised variance; labelling dominates higher up and it sits below. This is the crossing in the prior panel of Figure 4.

*Update speed.* The gain is set by how cleanly pay separates into permanent and transitory parts. More labelling shrinks the unlabelled residual variance  $V_u = \sigma_\psi^2 + 2(1 - \rho_g)\sigma_\theta^2$  relative to the total, so each squared surprise is a sharper signal about permanent risk and the steady-state Kalman gain is larger. The top, which labels most, updates fastest; the bottom, which labels least, filters the noisiest mix and updates slowest. Reading noise washes out of the gain, since it scales the squared signal and its sampling noise together. This is the  $\mu$  panel of Figure 4 (filter and algebra in Appendix F).

*Attribution share.* The agent updates her permanent belief only on the unlabelled part of a shock, but the regression has only the full shock to load on, so the labelled part adds movement to the regressor with no matching movement in the belief and attenuates the slope, as classical noise in a regressor would. Under joint Gaussianity of  $\psi$  and  $\theta$  (derivation in Appendix C),

$$\lambda_g = \frac{\sigma_\psi^2(\sigma_\psi^2 + 2(1 - \rho_g)\sigma_\theta^2)}{(\sigma_\psi^2 + 2\sigma_\theta^2)^2}, \quad (7)$$

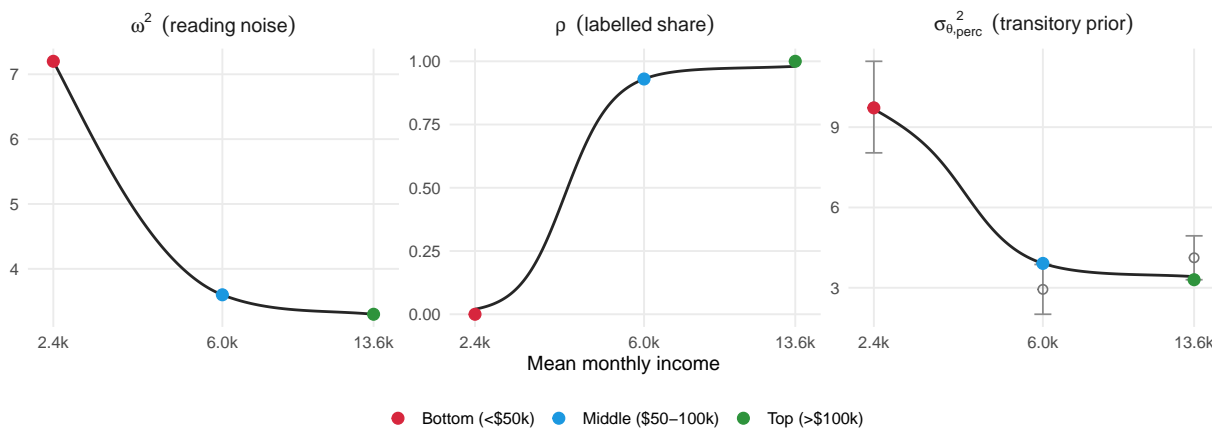
which decreases in  $\rho_g$  and equals the FIRE benchmark  $\lambda^*$  at  $\rho_g = 0$ . Through labelling alone, the bottom sits closest to that benchmark and the top furthest. Reading noise scales the slope down by a further factor, largest where the noise is largest (the bottom); the monotone  $\lambda$  column is the net of the two forces.

*A monotone reading-noise profile.* Reading noise enters only the prior and washes out of the gain, so the belief moments pin it solely through the prior level:  $\omega_g^2 = \sigma_{\theta,\text{perc},g}^2 - (1 - \rho_g)\sigma_\theta^2$  is a residual, and it inherits the non-monotonicity of the estimated prior. That non-monotonicity sits within the prior's sampling error, so I smooth it to the closest monotone-decreasing profile, one that holds the implied prior within one standard error at every group and matches the within-person hours-volatility gradient of Appendix E. Both pay-delivery primitives then move monotonically with income, the labelled share rising and the reading noise falling, and Figure 5 reports them,

---

<sup>4</sup> $\theta_t - \theta_{t-1}$  has variance  $2\sigma_\theta^2$  and first autocovariance  $-\sigma_\theta^2$ ;  $\nu_t - \nu_{t-1}$  has variance  $2\omega_g^2$  and first autocovariance  $-\omega_g^2$ . A transitory innovation of variance  $\omega_g^2$  is its exact observational twin.

and the perceived prior they imply, as continuous functions of income: a logistic profile for the share and a log-linear one for the noise, the links that keep a bounded share and a positive variance in range, each interpolated through the three groups, with the prior derived as  $(1 - \rho_g)\sigma_\theta^2 + \omega_g^2$ .



**Figure 5:** The two pay-delivery primitives and the perceived transitory prior they imply, by income. Points are the three income groups; lines interpolate them as continuous functions of income (the labelled share  $\rho$  through a logistic link, the reading noise  $\omega^2$  through a log link, and the prior derived as  $(1 - \rho)\sigma_\theta^2 + \omega^2$ ). Grey markers on the right panel are the GMM prior estimate  $\pm 1$  standard error; the calibrated prior lies within the band at every group. In  $(\%)^2$  except the dimensionless share.

### Information, not bias

Given  $(\omega_g^2, \rho_g)$ , the agent neither over- nor under-reacts to her signal relative to the Bayesian benchmark; her deviations from FIRE arise from what she can observe, not from how she reacts to it. The income gradient in beliefs is a gradient in information technology, not in rationality. The one departure from full rationality is the constant gain itself, a perpetual-learning assumption standard in adaptive-learning models, not a directional bias in how each signal is read. This deviation is what separates the mechanism from over-extrapolation accounts such as diagnostic expectations, which add a directional bias on top. Because the labelled share is observable as  $\rho_g$ , the mechanism embodies the superior-information critique of Kaufmann and Pistaferri (2009): a gap between perceived and econometric risk can reflect what the worker knows and the data cannot see, and here that knowledge is  $\rho_g$ .

## 9 A heterogeneous-agent model

To trace what those perceptions imply for behaviour, I compare each group under its estimated beliefs with an otherwise identical full-information benchmark in a standard heterogeneous-agent quantitative model. The gap between the two regimes, in precautionary wealth and the marginal propensity to consume, is the belief wedge.

### Setup

I embed the estimated belief structure in an Aiyagari-style buffer-stock economy, solved in partial equilibrium. Liquidity is group-specific and rises steeply with income: consistent with the HANK literature the bottom is near hand-to-mouth, and the top holds a multi-month buffer. This gradient

alone reproduces the steep empirical fall of MPCs with income, and the belief wedge operates on top of it.

Income follows the permanent-transitory process of Section 4,  $y_{i,t} = P_{i,t} \exp(\theta_{i,t})$  with  $P_{i,t} = P_{i,t-1} \exp(\psi_{i,t})$  and group-specific variances from Table 2. Workers do not move across groups: the perceived-risk object and the wage process condition on the same job, hours, and employer, so income evolves through within-job shocks alone, without the promotions or job switches that would carry a worker across the wide income bands. Following Carroll (2006), I normalise by permanent income: the state is cash-on-hand per unit of permanent income, the per-period draw is the transitory  $\exp(\theta_g)$ , and permanent shocks enter through the growth factor  $\exp(\psi_g)$ . The problem is scale-invariant, so I report each group's mean monthly SCE income (\$2,415 / \$6,014 / \$13,604) only to fix magnitudes.<sup>5</sup>

*Two regimes.* I solve each group twice. Under the *belief* regime the agent perceives wage risk through the estimated rule and her belief evolves; under *FIRE* she knows the true variances. Everything else is shared, so any gap is the belief wedge. FIRE here means correct wage-risk beliefs; job-loss risk is common to both regimes and nets out.

*Unemployment.* The SCE measures *same-job* risk, so the budget adds a standard two-state employment block: monthly separation 0.034 and job-finding 0.45 (Shimer, 2005), which imply a steady-state unemployment rate  $s/(s+f) = 7$  percent, with replacement income half of permanent. Because it is common across groups and regimes, it raises the level of precaution but adds nothing to the cross-group or belief-versus-FIRE comparison; beliefs update only while employed and freeze during unemployment.

### The household problem

$$V_g(a, y, \hat{\sigma}_\psi^2) = \max_c \left\{ u(c) + \beta \mathbb{E}[V_g(a', y', \hat{\sigma}_{\psi'}^2) \mid y, \hat{\sigma}_\psi^2] \right\}, \quad \text{s.t. } a' = Ra + y - c, \quad a' \geq 0.$$

The states are assets  $a$ , income  $y$ , and the permanent-variance belief  $\hat{\sigma}_\psi^2$ . The two information channels enter separately: the budget uses *realised* income, while expectations are taken under the agent's *perceived* process, so beliefs act only through the saving motive, and a higher  $\hat{\sigma}_\psi^2$  raises perceived risk and saving. The agent observes realised income each period but cannot decompose a given month's change into permanent and transitory parts, which is the signal-extraction problem of Section 8; for tractability the budget is written in permanent-income-normalised form using her running estimate of permanent income, so the normalisation is an approximation, not a claim that she observes the permanent component exactly. Beliefs then distort the perceived *variance* of future shocks, not the level of current income. The belief follows the rule of Section 5,  $\hat{\sigma}_{\psi,t}^2 = (1 - \mu_g)\hat{\sigma}_{\psi,t-1}^2 + \mu_g\lambda_g(\Delta\tilde{y}_t)^2$ , so it settles around  $\lambda_g \mathbb{E}[(\Delta\tilde{y})^2]$ , which the labelling discount holds below the realised permanent variance at the middle and top.

*Stationary equilibrium.* For each group and regime a stationary equilibrium is a value function, a policy  $c_g(a, y, \hat{\sigma}_\psi^2)$ , and an invariant distribution  $\Phi_g$  over  $(a, y, \hat{\sigma}_\psi^2)$  such that the policy solves the household problem at the fixed  $r$  and  $\Phi_g$  is invariant under the policy, the income process, the employment transitions, and the belief law of motion (degenerate at the truth under FIRE). I read

<sup>5</sup>Wage variances are stated in (%)<sup>2</sup>; in the income process they enter in fractional units (divided by 10<sup>4</sup>) before exponentiating, and the perceived-variance state is carried back in (%)<sup>2</sup>.

moments of  $\Phi_g$  off a simulated panel of  $N = 10,000$  households over 600 months (300 burn-in), with the belief and FIRE branches sharing common random numbers, so the belief-minus-FIRE gap is free of start-state and sampling noise.

## Calibration

The calibration fixes two objects per group: how patient it is ( $\beta_g$ ) and what it believes about wage risk ( $\mu_g$ ,  $\lambda_g$ , and the prior). The model is block-recursive: beliefs evolve on the realised income process and never on assets or consumption, so the belief parameters are identified by belief moments independently of  $\beta_g$ , and  $\beta_g$  is identified by wealth given beliefs. Efficient estimation therefore separates into two blocks, which I take in turn. Table 4 collects the calibration.

Parameter	Symbol	Value (Bot / Mid / Top)	Source or target
<i>A. Preferences and assets (common)</i>			
Risk aversion	$\gamma$	2	Assigned (standard)
Annual interest rate	$r$	2%	Assigned (standard)
Borrowing limit	$\underline{a}$	0	Assigned (standard)
Discount factor	$\beta_g$	0.80/0.93/0.95	Liquid wealth 0.5/2.0/3.0 mo. (SCE)
<i>B. Employment block (common across groups and regimes)</i>			
Separation rate	$s$	0.034/mo	Shimer (2005)
Job-finding rate	$f$	0.45/mo	Shimer (2005)
Replacement rate	$b$	0.5	Standard convention
<i>C. Realised income process by group, (%)<sup>2</sup></i>			
Permanent variance	$\sigma_\psi^2$	11.9/11.6/11.2	SIPP (Table 2)
Transitory variance	$\sigma_\theta^2$	2.5/4.4/5.8	SIPP (Table 2)
Mean monthly income	—	\$2,415/6,014/13,604	SCE (scale anchor only)
<i>D. Beliefs by group, belief regime (calibrated)</i>			
Update gain	$\mu_g$	0.628/0.912/0.950	Target: GMM (Table 6)
Attribution	$\lambda_g$	0.380/0.214/0.152	Forward-map under the prior (App F)
Transitory prior	$\sigma_{\theta, \text{perc}, g}^2$	9.72/3.91/3.30	$(1 - \rho_g)\sigma_\theta^2 + \omega_g^2$ , within GMM SE
Labelled share	$\rho_g$	0.00/0.93/1.00	Calibrated
Reading noise	$\omega_g^2$	7.20/3.60/3.30	Monotone profile (App F)

**Table 4:** Model calibration. Panels A–B are common; C is the realised SIPP process; D is the belief regime, calibrated to the Table 6 estimates (under FIRE, beliefs are held at the truth).

*Preferences, to match liquid wealth.* Preferences are CRRA ( $\gamma = 2$ ,  $r = 2$  percent annual). For each group I choose  $\beta_g$  so that its FIRE economy holds that group’s liquid wealth in the data, about 0.5/2.0/3.0 months (achieving 0.6/2.0/3.0; the bottom binds at an impatience floor), giving  $\beta = 0.80/0.93/0.95$ .<sup>6</sup> I solve by the endogenous grid method (Carroll, 2006) and simulate each regime.

*Beliefs, by indirect inference on the two primitives.* The three belief moments follow from the reading noise  $\omega_g^2$  and the labelled share  $\rho_g$ , plus a common Kalman-drift  $q$  that no group-level moment

<sup>6</sup>Targets from the SCE Household Finance module: non-retirement savings relative to income, the liquid-wealth concept of Kaplan et al. (2014). Median liquid wealth is about 0.2/1.5/3.2 months across groups, a steep rise with income; I round to 0.5/2.0/3.0.

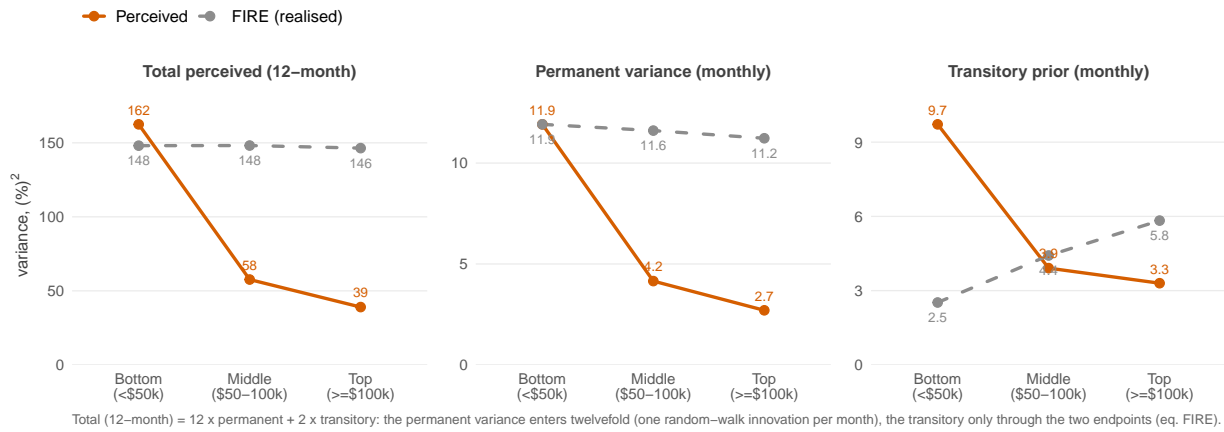
identifies. I match the update speed and the prior in levels and the attribution by simulating the cell-mean regressor and re-estimating (4) on the simulated panel, so the attenuation of Section 6 is generated rather than assumed and the attribution is matched in level. The simulated cell merge, with the common-shock share fixed at its measured value (about two percent), generates an attenuation of about five to one, so the structural attribution is of order 0.380/0.214/0.152 against a regression level of 0.077/0.046/0.026. The gain and the labelled share follow from the gain and attribution moments; the reading noise enters only the prior, where it is taken as the monotone profile of Section 8, holding the implied prior within the GMM standard errors and the implied attribution within the GMM band (Bot/Top and Mid/Top ratios 2.5/1.4 against the estimated 3.0/1.8, Figure 5).

Group	Gain $\mu_g$		Prior $\sigma_{\theta,\text{perc},g}^2$		Attrib. (level)	Primitives	
	data	fit	data	fit	$\lambda_g$ data / fit	$\rho_g$	$\omega_g^2$
Bot	0.628	0.627	9.75	9.72	0.077/0.075	0.00	7.20
Mid	0.912	0.907	2.94	3.91	0.046/0.046	0.93	3.60
Top	0.950	0.956	4.12	3.30	0.026/0.027	1.00	3.30

**Table 5:** Belief moments under the calibration. The gain matches the GMM in level and the labelled share  $\rho_g$  is set by the gain and the attribution ratios; the reading noise  $\omega_g^2$  is taken monotone (Section 8), so the implied prior lies within the GMM standard errors ( $z = -0.0/ + 1.0/ - 1.0$ ) rather than reproducing its point estimate. The attribution level is matched through the de-attenuation of Appendix F; the household block reads the structural attribution 0.380/0.214/0.152.

*Identification of the labelled share.* The fitted shares rise with income, as the mechanism and the performance-pay literature predict (Lemieux et al., 2009). I use a parametric bootstrap, in which I redraw the belief moments from their sampling distributions and refit, to obtain the sampling uncertainty of the calibrated shares (Appendix F): the bottom labels essentially nothing ( $\rho_{\text{Bot}}$  below 0.29 in 90 percent of draws). In this bootstrap, the middle and top shares are identified only as an interval ( $\rho_{\text{Mid}}$  in [0.53, 1.00],  $\rho_{\text{Top}}$  in [0.69, 1.00]): the attribution ratio that pins the spread is imprecise and sits near the largest value labelling alone can produce, so the fit is pushed to the bound. The gradient is robust, and the corner values 0.93 and 1.00 are an upper region and not necessarily fixed point estimates.

*What the calibrated beliefs imply.* The calibration pins what each group perceives before any behaviour. Figure 6, reported numerically in Table 9 (Appendix G), sets the implied perceived variance against FIRE, in the same SIPP units, so the  $\sigma_\psi^2$  inflation that blocks an absolute comparison with the SCE (Section 4) cancels. Perceived risk falls steeply with income, from 162 to 39, while FIRE is flat near 148. The decomposition shows the fall is the permanent labelling discount: the permanent belief, which drives saving, equals the realised value at the bottom ( $\rho = 0$ ) and is discounted to 36 and 24 percent of it at the middle and top. The bottom's perceived total sits just above FIRE, but that is its inflated transitory prior alone, not an over-perception of the permanent risk that matters for saving.

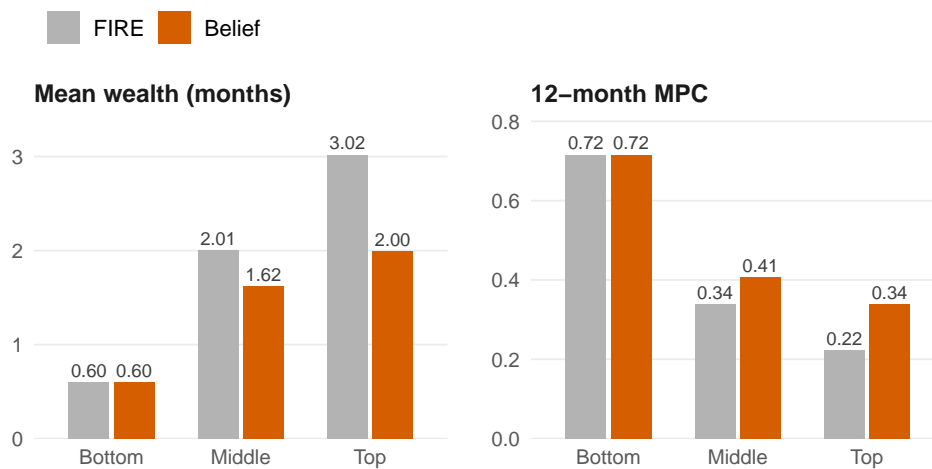


**Figure 6:** Calibrated perceived variance against the FIRE benchmark, by income group, in  $(\%)^2$ . Left: total at the 12-month horizon. Centre and right: the monthly permanent-variance belief and transitory prior, each against its realised value. The total is  $12\hat{\sigma}_\psi^2 + 2\sigma_{\theta,perc}^2$ .

## 10 Partial-equilibrium results

### Steady state

Figure 7 compares the belief regime with an otherwise identical FIRE benchmark; the underlying numbers are in Table 10 (Appendix G).



**Figure 7:** Mean liquid wealth (months of permanent income) and the 12-month MPC by income group, FIRE versus the belief regime (Table 10).

Liquidity sets the slope and beliefs shift the level. The cross-group MPC gradient is a liquidity object: under FIRE the 12-month MPC already falls from 0.72 at the bottom to 0.34 and 0.22 above, tracking the liquid-wealth gradient. Beliefs add a within-group wedge on top, +0.0, +6.8, and +11.5 points, zero at the constrained bottom and rising with income. The two channels load on opposite margins: liquidity generates the dispersion *across* groups, beliefs raise the level *within* a group, and only where the household is unconstrained enough to act.

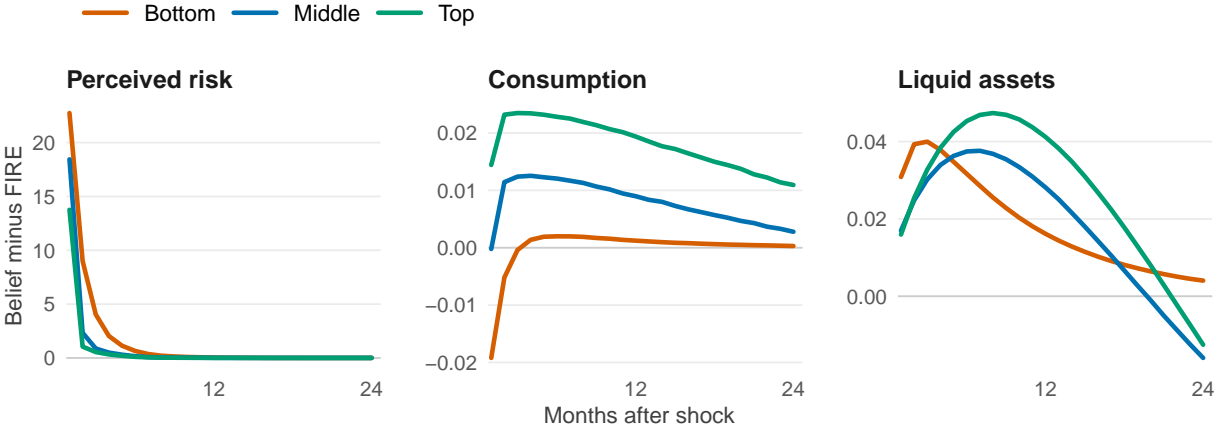
Beliefs barely move the bottom (within one percent on wealth and effectively nothing on the MPC): its perceived risk is the highest of the three, but the distortion is entirely in the transitory prior,

and with nothing labelled its permanent belief is unbiased. Being near hand-to-mouth, it has neither the room nor the wedge to act on. The middle and top respond sharply, holding 19 and 34 percent less wealth and consuming 7 and 12 points more of a windfall: by labelling transitory pay away they under-perceive permanent risk and save less. The wedge falls on the groups that distort permanent risk, the component that drives saving, not on the group that distorts only the transitory prior. The labelled belief lies below the truth at all times, so the wedge is a steady-state object, not merely a response to shocks.

**Dynamics**

With prices fixed, the experiment traces the belief channel through saving alone. I begin with the shock one would expect to matter most, a first-order recession, and find it barely moves the wedge; the belief channel is driven instead by a shock to risk itself.

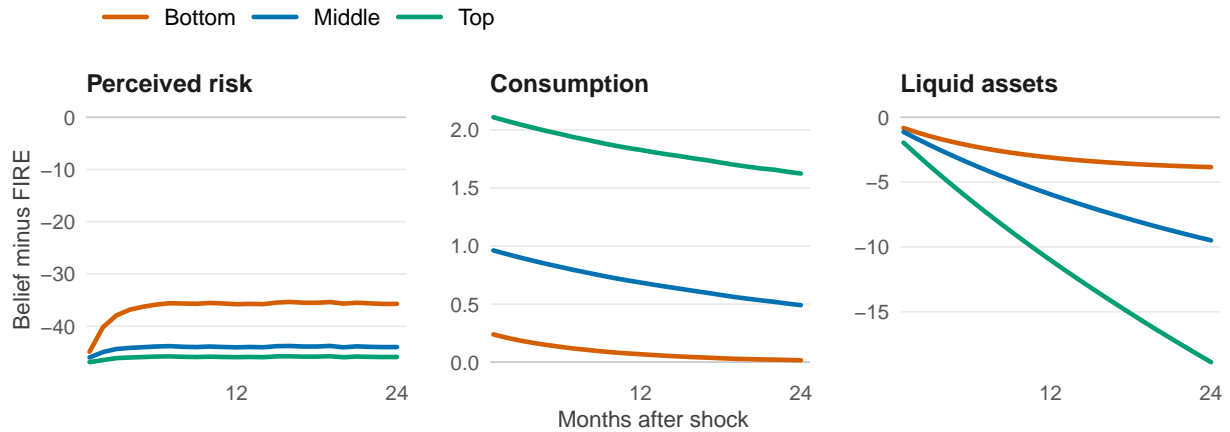
*A first-order shock.* A single negative innovation to the common permanent component, a permanent 3 percent level fall of the kind a recession delivers, traced against each regime’s own no-shock path (Figure 8). Consumption falls about 2.5 percent on impact in both regimes, and perceived risk rises 13–23 (%)<sup>2</sup> through the squared signal, but the belief regime adds at most a few hundredths of a percent to the consumption response, fading within a quarter. A one-off level shock is a transient spike in the squared signal, and any constant-gain rule discounts it geometrically, so it barely moves the belief stock. The wedge lives in the chronic level of perceived risk, not in its reaction to one shock, meaning it is most affected when risk changes, not when the level does.



**Figure 8:** Response to a single 3 percent permanent (recession) income fall, belief regime minus FIRE, by income group, on independent scales. Perceived risk in (%)<sup>2</sup>; consumption and liquid assets in percent of the no-shock baseline.

*A rise in income risk.* A persistent +40% rise in the realised permanent variance. FIRE recognises it at once and rebuilds its buffer. The belief regime records a share  $\lambda$  of the now-larger squared surprises as permanent, so its perceived permanent variance rises by only  $\lambda$  of the truth, about 0.38/0.21/0.15 of the FIRE rise (Figure 9, left). It under-reacts, most at the top where  $\lambda$  is lowest, and under-saves: it holds 4/9/19 percent less liquid wealth than FIRE after two years and is still falling behind, while its consumption sits +0.2/ +1.0/ +2.1 percent above FIRE on impact. The wedge is largest at the top for two reasons: it is unconstrained, so the appropriate response is

large, and its labelling discount is largest, so it perceives the least of the rise.



**Figure 9:** Response to a persistent 40% rise in realised permanent wage variance, belief regime minus FIRE, by income group, on independent scales. Perceived risk is the gap in the rise in twelve-month variance, in  $(\%)^2$ ; consumption and liquid assets are in percent of the no-shock baseline.

The belief distortion is largest at the bottom, but the behavioural wedge falls on the top: liquidity selects which distortion reaches behaviour. The bottom, whose beliefs are most distorted, is near hand-to-mouth and cannot act on them; the top, which under-perceives the permanent risk that drives saving, is cushioned and does. This is the dynamic counterpart of the steady-state cross-section. In general equilibrium these wedges would move  $r$  by shifting aggregate saving, with the middle and top dominating the response; a credible magnitude requires a two-asset closure and is left to future work.

## References

- Blundell, R. and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1):115–143.
- Bound, J. and Krueger, A. B. (1991). The extent of measurement error in longitudinal earnings data: Do two wrongs make a right? *Journal of Labor Economics*, 9(1):1–24.
- Broer, T., Kohlhas, A., Mitman, K., and Schlafmann, K. (2022). Expectation and wealth heterogeneity in the macroeconomy. Working paper.
- Caplin, A., Gregory, V., Lee, E., Leth-Petersen, S., and Sæverud, J. (2023). Subjective earnings risk. NBER Working Paper 31019, National Bureau of Economic Research.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.
- Grigsby, J., Hurst, E., and Yildirmaz, A. (2021). Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review*, 111(2):428–471.
- Guerreiro, J. (2023). Belief disagreement and business cycles. Job Market Paper, Northwestern University.
- Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of US workers reveal about lifecycle earnings dynamics? *Quarterly Journal of Economics*, 136(4):2349–2391.
- Halvorsen, E., Holter, H. A., Ozkan, S., and Storesletten, K. (2024). Dissecting idiosyncratic earnings risk. *Review of Economic Studies*.
- Hardy, B. and Ziliak, J. P. (2014). Decomposing trends in income volatility: The “Wild Ride” at the top and bottom. *Economic Inquiry*, 52(1):459–476.
- Hoynes, H., Miller, D. L., and Schaller, J. (2012). Who suffers during recessions? *Journal of Economic Perspectives*, 26(3):27–48.
- Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, pages 77–138.
- Kaufmann, K. and Pistaferri, L. (2009). Disentangling insurance and information in intertemporal consumption choices. *American Economic Review: Papers & Proceedings*, 99(2):387–392.
- Kim, G. and Binder, C. (2023). Learning-through-survey in inflation expectations. *American Economic Journal: Macroeconomics*, 15(2):254–278.
- Kuchler, T. and Zafar, B. (2019). Personal experiences and expectations about aggregate outcomes. *Journal of Finance*, 74(5):2491–2542.
- Lemieux, T., MacLeod, W. B., and Parent, D. (2009). Performance pay and wage inequality. *Quarterly Journal of Economics*, 124(1):1–49.
- Ma, Y., Ropele, T., Sraer, D., and Thesmar, D. (2020). A quantitative analysis of distortions in managerial forecasts. NBER Working Paper 26830, National Bureau of Economic Research.
- Meghir, C. and Pistaferri, L. (2004). Income variance dynamics and heterogeneity. *Econometrica*,

72(1):1–32.

Morduch, J. and Schneider, R. (2017). *The Financial Diaries: How American Families Cope in a World of Uncertainty*. Princeton University Press.

Pischke, J.-S. (1995). Individual income, incomplete information, and aggregate consumption. *Econometrica*, 63(4):805–840.

Rozsypal, F. and Schlafmann, K. (2023). Overpersistence bias in individual income expectations and its aggregate implications. *American Economic Journal: Macroeconomics*, 15(4):331–371.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49.

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

Wang, T. (2023). Perceived versus calibrated income risks in heterogeneous-agent consumption models. Bank of Canada Staff Working Paper 23-59.

## A From EWMA to observable AR(1)

Solving the level identity (3) for the belief gives  $\hat{\sigma}_{\psi,t}^2 = (\bar{\sigma}_{\text{perc},t}^2 - 2\sigma_{\theta,\text{perc}}^2)/12$ . Substituting both belief terms into (2) and multiplying by 12,

$$\bar{\sigma}_{\text{perc},t}^2 - 2\sigma_{\theta,\text{perc}}^2 = (1 - \mu)[\bar{\sigma}_{\text{perc},t-1}^2 - 2\sigma_{\theta,\text{perc}}^2] + 12\mu\lambda(\Delta e_{t-1})^2.$$

Collecting the intercept ( $2\mu\sigma_{\theta,\text{perc}}^2$ ), adding group subscripts, a person fixed effect  $\alpha_i$ , and a shock  $u_{i,t}$  gives (4).

## B 12-month FIRE variance

Rolling  $e_t = p_t + \theta_t$  forward  $H$  months, the permanent piece accumulates one fresh innovation per month,  $\text{Var}(p_{t+H} - p_t) = H\sigma_{\psi}^2$ , while the iid transitory piece contributes only its endpoints,  $\text{Var}(\theta_{t+H} - \theta_t) = 2\sigma_{\theta}^2$ . The shocks are independent, so  $\text{Var}(e_{t+H} - e_t) = H\sigma_{\psi}^2 + 2\sigma_{\theta}^2$ , which at  $H = 12$  is (1).

## C Closed form for $\lambda_g$ under partial labelling

Split the transitory shock  $\theta_t = \theta_t^{\text{lab}} + \theta_t^{\text{unlab}}$  into independent pieces with variances  $\rho_g\sigma_{\theta}^2$  and  $(1 - \rho_g)\sigma_{\theta}^2$ . Then  $\Delta e_t = u_t + l_t$ , with unlabelled residual  $u_t = \psi_t + \theta_t^{\text{unlab}} - \theta_{t-1}^{\text{unlab}}$  and labelled component  $l_t = \theta_t^{\text{lab}} - \theta_{t-1}^{\text{lab}}$ , so  $V_u = \sigma_{\psi}^2 + 2(1 - \rho_g)\sigma_{\theta}^2$ ,  $V_l = 2\rho_g\sigma_{\theta}^2$ , and  $V_u + V_l = \sigma_{\psi}^2 + 2\sigma_{\theta}^2$ .

The belief responds only to  $u_t^2$ . Steady-state Bayesian consistency  $\mathbb{E}[\hat{\sigma}_{\psi}^2] = \sigma_{\psi}^2$  pins the belief-to- $u^2$  proportionality at  $\sigma_{\psi}^2/V_u$ , so the slope identified by (4) is  $\lambda_g = (\sigma_{\psi}^2/V_u) \text{Cov}(u^2, (u+l)^2)/\text{Var}((u+l)^2)$ . Under joint Gaussianity with  $u \perp l$ ,  $\text{Cov}(u^2, (u+l)^2) = 2V_u^2$  and  $\text{Var}((u+l)^2) = 2(V_u + V_l)^2$ , so

$$\lambda_g = \frac{\sigma_{\psi}^2 V_u}{(V_u + V_l)^2} = \frac{\sigma_{\psi}^2 (\sigma_{\psi}^2 + 2(1 - \rho_g)\sigma_{\theta}^2)}{(\sigma_{\psi}^2 + 2\sigma_{\theta}^2)^2},$$

which is (7). When reading noise (5) is included, the agent updates on  $(u_t + \Delta\nu_t)^2$  while the regressor is unchanged, and the slope is scaled by  $V_u/(V_u + 2\omega_g^2) < 1$ , strongest where  $\omega^2$  is largest. The calibration of Section 9 disciplines the trade-off.

## D Belief estimates and robustness

Table 6 reports the headline joint-GMM estimates of (4) plotted in Figure 4. The rest of this appendix checks their robustness: D.1 adds an occupation  $\times$  time fixed effect to the shock construction, D.2 varies the instrument lags, the trim, and the COVID handling one at a time, and D.3 reports the sign-split shock.

Group	$\sigma_{\theta}^2$ (SIPP truth)	$\sigma_{\theta, \text{perc}, g}^2$ (GMM)	$\lambda_g$ (GMM)	$\mu_g$ (GMM)	Hansen $J$ $p$	$N$
Bot	2.52	<b>9.75</b> (1.71)	+ <b>0.077</b> (0.030)	0.628 (0.055)	0.20	1,531
Mid	4.42	2.94 (0.93)	+0.046 (0.015)	0.912 (0.089)	0.23	7,501
Top	5.83	4.12 (0.82)	+0.026 (0.013)	0.950 (0.060)	0.24	11,331
Wald ( $B = T$ )	—	$p = \mathbf{0.003}$	$p = 0.12$	$p = \mathbf{0.0001}$	—	—

**Table 6:** Joint Blundell-Bond system GMM on (4), by income group (the estimates of Figure 4). Two-way clustered SEs (respondent, quarter) in parentheses. Instruments:  $\bar{\sigma}_{\text{perc}}^2$  at lags 3–4, the shock at lags 1–3, plus BB level moments. The Hansen  $J$  never rejects.

### D.1 Occupation $\times$ time fixed effect

Adding an occupation  $\times$  year-month fixed effect to the wage-shock construction (the interaction is the active term, since an additive occupation effect cancels in the within-spell difference) preserves the persistence and attribution gradients and sharpens the attribution; the prior gradient holds in direction but loses cross-group significance, so part of the bottom’s inflated prior reflects cross-occupation transitory dispersion. The labelling channel behind  $\lambda$  is robust to occupation controls; the reading-noise channel behind the prior operates partly through cross-occupation variation.

Group	Canonical			Occupation $\times$ time FE		
	$\hat{\mu}_g$	$\hat{\lambda}_g$	$\hat{\sigma}_{\theta, \text{perc}, g}^2$	$\hat{\mu}_g$	$\hat{\lambda}_g$	$\hat{\sigma}_{\theta, \text{perc}, g}^2$
Bot	0.628 (0.055)	+ <b>0.077</b> (0.030)	<b>9.75</b> (1.71)	0.649 (0.063)	+ <b>0.091</b> (0.039)	<b>7.20</b> (2.66)
Mid	0.912 (0.089)	+0.046 (0.015)	2.94 (0.93)	0.924 (0.091)	+0.036 (0.010)	2.74 (0.96)
Top	0.950 (0.060)	+0.026 (0.013)	4.12 (0.82)	0.957 (0.061)	+0.017 (0.008)	4.17 (0.74)
Wald $\mu$ (T–B)		$p = \mathbf{0.0001}$			$p = \mathbf{0.0004}$	
Wald $\lambda$ (B–T)		$p = 0.118$			$p = 0.062$	
Wald $\sigma^2$ (B–T)		$p = \mathbf{0.003}$			$p = 0.274$	

**Table 7:** Canonical specification versus the same regression with an occupation  $\times$  year-month fixed effect in the wage-shock construction. Standard errors in parentheses.

### D.2 Lag depth, trim, and COVID exclusion

Each row of Table 8 varies one element of the canonical specification. The persistence and attribution gradients are robust across the lag variants, with the attribution Wald tighter than canonical in three of four. The within-group trim preserves the qualitative pattern but loosens both Walds. Excluding COVID sharpens the attribution gradient and breaks the prior gradient’s significance, so part of the bottom’s elevated prior is identified from the COVID year. The Hansen  $J$  never rejects.

Variant	$\hat{\mu}_g$			$\hat{\lambda}_g$			$\hat{\sigma}_{\theta, \text{perc}, g}^2$			Wald $p$ (B vs T)		
	B	M	T	B	M	T	B	M	T	$\mu$	$\lambda$	$\sigma^2$
Canonical	0.628	0.912	0.950	+0.077	+0.046	+0.026	9.75	2.94	4.12	<b>0.0001</b>	0.118	<b>0.003</b>
Trim 2.5/97.5	0.816	0.847	0.897	+0.059	+0.053	+0.014	11.23	4.07	4.86	0.110	0.352	<b>0.000</b>
$\bar{\sigma}^2$ lags 2-3	0.648	0.951	0.876	+0.075	+0.050	+0.020	9.96	2.85	4.44	<b>0.000</b>	0.079	<b>0.003</b>
$\bar{\sigma}^2$ lags 4-5	0.588	0.901	0.948	+0.102	+0.046	+0.021	9.03	3.00	4.38	<b>0.000</b>	0.061	<b>0.032</b>
Shock lags 0-1	0.683	0.937	0.940	+0.082	+0.016	+0.012	9.76	4.41	4.90	<b>0.015</b>	0.063	<b>0.010</b>
Shock lags 2-4	0.675	0.908	0.948	+0.095	+0.040	+0.027	8.76	3.20	4.10	<b>0.003</b>	<b>0.008</b>	<b>0.011</b>
Exclude 2020-21	0.636	0.995	0.886	+0.139	+0.043	+0.013	6.50	3.12	4.88	<b>0.017</b>	0.055	0.612


**Table 8:** One-at-a-time variations around the canonical specification. Bold  $p$ -values reject cross-group equality at the 5% level.

### D.3 Asymmetric extrapolation

Pooling across groups and splitting the shock by sign, both components load positively, with positive news carrying roughly twice the weight of negative news ( $\lambda^+ = 0.055$  versus  $\lambda^- = 0.027$ ), significant across cluster choices ( $N = 20,363$ , Hansen  $J$   $p = 0.11$ ). Upside news is read as informative about future earning power; downside news is more often dismissed as transitory, the opposite of loss aversion.

## E Hours by income group

Both primitives operate through hours: fragmented, variable hours generate the reading noise  $\omega_g^2$ , and transitory shocks delivered through hours are of the unlabelled kind and lower  $\rho_g$ . Within-person hours volatility  $\mathbb{E}[(\Delta \log h)^2]$  falls monotonically with income (4.10/3.06/2.36 in  $(\%)^2$ ), consistent with reading noise being highest at the bottom (Figure 10). The hours proxy is a one-directional check on  $\omega^2$ : it supports a reading-noise channel that falls with income, which is the monotone profile adopted in the calibration (Section 8), while the prior's level is supplied by the labelling channel higher up (Appendix F).



**Figure 10:** Mean weekly hours (left) and within-person hours volatility  $\mathbb{E}[(\Delta \log h)^2 | g, t]$  in  $(\%)^2$  (right) by income group. SIPP matched-stayer panel, quarterly with 4-quarter trailing MA.

## F The filter and the identification of the belief block

This appendix gives the filter behind the gain, shows that both primitives are necessary, reports how the attribution level is identified, and reports the sampling uncertainty of the labelled share.

*The filter.* The agent’s monthly signal is one squared pay reading,  $z_t = (\Delta \tilde{y}_t)^2$ , with mean  $\sigma_\psi^2 + V_{x,g}$  and, under Gaussianity, variance  $2(\sigma_\psi^2 + V_{x,g})^2$ , where  $V_{x,g} = 2(1 - \rho_g)\sigma_\theta^2 + 2\omega_g^2 = 2\sigma_{\theta,\text{perc},g}^2$ . Treating the permanent variance as a random walk with drift  $q$  observed through the attributed signal  $\lambda_g z_t$ , the steady-state Kalman gain is

$$\mu_g = \frac{\bar{P}_g}{\bar{P}_g + R_g}, \quad R_g = \lambda_g^2 2(\sigma_\psi^2 + V_{x,g})^2, \quad \bar{P}_g = \frac{1}{2}(q + \sqrt{q^2 + 4qR_g}).$$

Substituting  $\lambda_g$  collapses the observation noise to  $R_g = 2(\sigma_\psi^2)^2(V_u/S)^4$ , so the gain depends on the labelled share and not on reading noise, which cancels out of the signal-to-noise ratio.

*Both primitives are necessary.* The perceived prior crosses realised transitory variance (+7.2 at the bottom,  $-1.5$  and  $-1.7$  at the middle and top). Reading noise can only raise the prior above  $\sigma_\theta^2$  and labelling can only lower it below, so the crossing cannot be reproduced by either alone. Refitting under each restriction confirms this: imposing  $\rho_g \equiv 0$  leaves the middle and top priors stuck at or above their realised floor and flattens the gain gradient ( $J = 24.7$ , 4 df,  $p < 0.001$ ); imposing  $\omega_g^2 \equiv 0$  leaves the bottom prior stuck at its realised floor of 2.52 against a target of 9.75 ( $J = 24.9$ , 4 df,  $p < 0.001$ ). The two-primitive fit is not rejected ( $J = 0.03$ ).

*Attribution level and the attenuation.* The cell-mean regressor attenuates the attribution through finite-cell sampling noise, not through a classical proxy error: the cell mean tracks only the common component of a worker’s shock, and the slope is attenuated by the share of the squared-shock variance that is common. The between-cell variance share of  $(\Delta e)^2$  in SIPP is about two percent. Simulating the SCE/SIPP cell merge at that share and re-estimating (4) reproduces an empirical attenuation of about five to one, so a structural attribution of order 0.380/0.214/0.152 is recovered in level and matched to the regression level 0.077/0.046/0.026. The level is therefore an output of the measured common-shock share, not an assumption.

*Sampling uncertainty of the labelled share.* The intervals reported with the labelled share in Section 9 come from a parametric bootstrap, since the share is calibrated by indirect inference rather than estimated as a regression coefficient and so has no standard error of its own. In each of 200 replications I draw the gain, prior, and attribution for the three groups from independent normals centred at their GMM estimates with their two-way-clustered standard errors, refit  $(\omega_g^2, \rho_g, q)$ , and retain the implied  $\rho_g$ . The 5th and 95th percentiles across replications are [0.00, 0.29] for  $\rho_{\text{Bot}}$ , [0.53, 1.00] for  $\rho_{\text{Mid}}$ , and [0.69, 1.00] for  $\rho_{\text{Top}}$ . The middle and top reach the upper bound because the attribution ratio that pins the spread is imprecise (Wald  $p = 0.12$ ), so those corners are not point-identified. The prior moment, absorbed by  $\omega_g^2$ , contributes nothing to the over-identifying  $J$ ; the over-identifying content is the gain and the attribution.

*The reading-noise profile.* Because  $\omega_g^2$  matches the prior one-for-one and is gain-invariant, the belief moments identify it only through the prior level; the U-shape of the residual (a middle of 2.64 below a top of 4.14) is not a separate finding. For the model I adopt the closest monotone-decreasing profile,  $\omega^2 = 7.2/3.6/3.3$ , the projection of the residual onto the decreasing sequences in the prior’s standard-error metric. It holds the implied prior within the estimated bands ( $z = -0.0/+1.0/-1.0$ ), follows the hours-volatility gradient of Appendix E, and leaves the gain unchanged. It shifts the implied attribution within the GMM band (Bot/Top and Mid/Top ratios 2.5/1.4 against the

estimated 3.0/1.8), and the permanent belief  $\sigma_\psi^2 (V_u/S)^2$  is itself reading-noise-invariant, so the steady-state wedge of Section 10 is unchanged.

*Why three groups.* The continuous profiles of Figure 5 interpolate three income groups; finer bins cannot anchor them. The SCE elicits income in brackets whose extremes are thin, and splitting the terciles into the five native brackets and re-estimating breaks the belief block exactly where it is split: the lowest bracket (under \$20k, 102 respondents) retains no cell that clears the estimation floor, the attribution turns negative in the lower-middle bracket, and the gain exceeds one (an inadmissible negative persistence) in the top two brackets; only the middle bracket, which is the unchanged tercile, survives. System GMM with lagged instruments and two-way clustering needs hundreds of within-person spells per cell, and the bracketed extremes do not supply them. The terciles are the finest split the survey powers, and the continuum is an interpolation through them, not an estimate.

*Link functions.* The income profiles in Figure 5 (Section 8) interpolate the three groups in the natural scale of each object: the labelled share  $\rho \in [0, 1]$  through a logit link, inverting through the logistic, and the reading noise  $\omega^2 > 0$  through a log link. Each keeps the interpolant in range where a common link would not, since a log scale leaves the share unbounded above one and a logit scale is undefined for a variance above one. With three monotone anchors the choice does not move the calibrated group values; it sets only the shape between groups and any extrapolation.

## G Steady-state model output

Table 9 gives the numbers behind Figure 6: the calibrated perceived variance against the full-information benchmark, by income group.

Group	$\hat{\sigma}_\psi^2$	Perceived $\bar{\sigma}_{\text{perc}}^2$	FIRE	Ratio
Bot	11.9	162	148	1.10
Mid	4.2	58	148	0.39
Top	2.7	39	146	0.27

**Table 9:** Steady-state perceived variance against the FIRE benchmark, in (%)<sup>2</sup> at the 12-month horizon.  $\hat{\sigma}_\psi^2 = \sigma_\psi^2 (V_u/S)^2$ ;  $\bar{\sigma}_{\text{perc}}^2 = 12\hat{\sigma}_\psi^2 + 2\sigma_{\theta, \text{perc}}^2$ ; FIRE =  $12\sigma_\psi^2 + 2\sigma_\theta^2$ .

Table 10 gives the numbers behind Figure 7: mean liquid wealth and the 12-month MPC, belief regime versus FIRE.

Group	Mean wealth (months)			12-month MPC		
	FIRE	Belief	$\Delta$	FIRE	Belief	$\Delta$
Bot	0.60	0.60	-0%	0.72	0.72	+0.0pp
Mid	2.01	1.62	-19%	0.34	0.41	+6.8pp
Top	3.02	2.00	-34%	0.22	0.34	+11.5pp

**Table 10:** Belief regime versus FIRE, by income group, under group-specific liquidity. Wealth in months of permanent income; MPC over 12 months out of a small transitory windfall. PE,  $r = 2$  percent,  $\beta = 0.80/0.93/0.95$ .