The Legibility of Income Risk*

Sven van Holten Charria†

Abstract

Do households perceive the income risk that macroeconomic models assign them? Job stayers in households below $50k report about 1.8 times the wage uncertainty of those above $100k, hold these beliefs more persistently, and revise them more strongly following realised wage shocks, while the realised risk estimated from matched same-job spells shows no comparable gradient. Under an estimated learning rule, the bottom group over-perceives its transitory risk, while the middle and top groups under-perceive their permanent risk. I trace both patterns to how legibly pay separates into these two components: the labelled share of transitory income rises from about one-eighth at the bottom to over one-third at the top. In a life-cycle economy, these beliefs reduce liquid wealth by five percent at the bottom and fourteen percent at the top, where households consume over two percentage points more of a windfall within the year. Equating perceived with realised risk therefore overstates precautionary wealth and understates consumption responses, particularly among high-income households.

1 Introduction

How much risk does a worker believe her own wage carries? In the heterogeneous-agent models now standard in macroeconomics, that belief is a primitive of behaviour: an agent who perceives more risk to her future earnings holds a larger precautionary buffer and spends less out of a windfall, whether through prudence or an occasionally binding borrowing constraint. The level and heterogeneity of income risk therefore shape the distribution of wealth, the distribution of marginal propensities to consume, and, through both, the transmission of aggregate shocks and policy. Yet these models rarely measure perceived income risk, and instead simply assign it.

The standard practice is to calibrate income risk indirectly. The econometrician estimates the dispersion of realised earnings in panel data, decomposes it into permanent and transitory components (Meghir and Pistaferri, 2004), and assigns the resulting variances to the agent as the risk she perceives. Perceived risk is thus equated with realised risk: the agent is assumed to know the variances the econometrician recovers and nothing beyond them. The assumption is convenient but strong, and data on subjective expectations now allow it to be tested.

I start by showing that this assumption misses a systematic income gradient in perceived wage risk. The New York Fed's Survey of Consumer Expectations (SCE) elicits from each respondent a density forecast of her own wage growth over the next twelve months, conditional on remaining

*I am grateful to Ricardo Reis, Jonathon Hazell, Marco Bellifemine, and participants of the LSE Pre-Doctoral seminars for their valuable comments and feedback, and I am especially indebted to my fellow pre-docs Nicholas Tokay and James Symons-Hicks for their willingness to keep on listening to my endless ideas.

†Pre-doctoral Research Assistant at the Centre for Macroeconomics, London School of Economics. E-mail: s.m.van-holten-charria@lse.ac.uk.

in the same job, with the same hours and employer. Among such workers, those in households earning under $50k report wage uncertainty about 1.8 times that of workers in households earning above $100k. They also hold that belief more persistently within person and revise it more strongly following a realised wage shock. None of these gradients appears in the realised wage process: using matched same-job spells in the Survey of Income and Program Participation (SIPP), I find little corresponding income gradient in realised risk, if anything a mild increase with income. The discrepancy is therefore not simply a matter of different groups facing different amounts of wage risk.

The gap becomes clearer when realised risk is decomposed into permanent and transitory components. I find that the bottom group over-perceives its transitory risk, while the middle and top groups under-perceive their permanent risk. Thus, groups lie on opposite sides of their own full-information benchmarks, and total perceived risk crosses the benchmark as well. This two-sided pattern is difficult to reconcile with a single directional source of belief error. I argue instead that it reflects how legibly pay separates into its permanent and transitory components. Higher in the distribution, transitory pay is more often explicitly labelled, for example, as a commission or bonus alongside a regular salary. Workers can therefore net out temporary payments before forming beliefs about persistent earning power, pushing perceived risk downward. At the bottom, pay is more fragmented across variable hours, tips, and irregular supplements. When these components are harder to identify, the resulting measurement noise pushes perceived risk upward.z

In Section 4, I formalise this idea as a model of belief formation. In the spirit of Pischke (1995), the worker faces a signal-extraction problem: each pay change mixes permanent and transitory shocks, and she must infer the variance of permanent earning power from this noisy signal. Her belief is an exponentially weighted average of past squared surprises, governed by two primitives: the labelled share of pay and the noise with which the remainder is read. She assigns a fixed share of each surprise to the permanent component and does not increase that share when labelling makes the signal cleaner.

The mechanism yields four results. First, greater labelling reduces the weight assigned to each surprise and therefore lowers perceived permanent risk. Second, the perceived transitory prior equals the unlabelled component plus reading noise, so it can lie either above or below its realised counterpart. Third, total perceived risk can consequently lie on either side of realised risk, depending on how legibly pay is delivered. This provides a mechanism for an income gradient in beliefs without an income gradient in the underlying wage process. Fourth, the behavioural consequences depend on which component is misperceived: permanent risk affects precautionary saving, whereas transitory risk has a much weaker saving consequence. The resulting distortion can therefore be largest among high-income workers even though their total perceived risk is the lowest.

In Section 5, I estimate the updating rule within person by income group and find that it organises the three empirical facts. The bottom group updates most slowly, assigning about three times as much of each squared surprise to the permanent component as the top group. Its high level of perceived risk instead comes from a transitory prior almost four times its realised counterpart, while the middle and top groups' transitory priors are at or below their realised levels. I then measure the mechanism's key primitive, the labelled share of within-spell pay movements, using the same panel that supplies the realised wage process. The labelled share rises from about one-eighth at the bottom to more than one-third at the top. Feeding these measured shares into the model predicts the middle and top groups' prior levels to within a standard error, reproduces the ordering of updating gains, and accounts for roughly half of the estimated income gradient in the response to news. The bottom group's prior remains substantially above the level implied by its labelled share; I interpret the residual as the reading noise generated by fragmented pay. The evidence therefore points to an income gradient in how pay is delivered, rather than in how risky pay is.

Section 6 embeds these measured beliefs in a life-cycle buffer-stock economy in which workers move across income groups as permanent earning power accumulates. Conditional on their income group, workers differ only in the risks they face, namely the realised wage variances and the job-loss and job-finding rates, while beliefs differ according to the estimates above. Relative to an otherwise identical full-information calibration, the measured beliefs reduce liquid wealth by five percent at the bottom and fourteen percent at the top. The MPC effect is concentrated at the top: these workers hold the most downward-biased beliefs about permanent risk and consume more than two percentage points more of a windfall within the year. The gap in MPCs between the bottom and top therefore narrows by about one-sixth. The bottom group, by contrast, combines a large transitory-risk overestimate with an approximately unbiased permanent-risk belief, leaving its saving behaviour little changed. In one-asset general equilibrium, the belief economy supplies modestly less capital and generates a slightly higher interest rate. Aggregate effects are therefore much smaller than the cross-sectional effects.

The same beliefs also govern the economy's response to changes in risk and pay structure. A reform that raises the labelled share of low-paid work to the level observed at the top reduces precautionary buffers in the newly legible groups by about three percent. Following a persistent rise in permanent risk, the belief economy rebuilds fifteen to nineteen percent less precautionary buffer than the full-information economy over two years because households incorporate only part of the increase into their perceived permanent risk. A full-information calibration therefore overstates the precautionary response to rising permanent risk. By contrast, a recession that changes neither the delivery of pay nor the underlying risk process leaves the two regimes with the same buffer response.

The paper thus makes three contributions. It documents an income gradient in perceived wage risk that is absent from the realised wage process; it shows that this gradient extends to the persistence and updating of beliefs and can place different income groups on opposite sides of their own full-information risk benchmarks; and it proposes and measures a mechanism based on the legibility of pay, with quantitatively important consequences for precautionary saving and the distribution of MPCs.

1.1 Related Literature

The closest paper is Wang (2023), who use the SCE's density forecasts to show that perceived wage risk lies well below the risk conventionally calibrated from panel data, including within demographic groups. They further show that feeding this lower and more heterogeneous risk into an incomplete-market model raises wealth inequality and the share of hand-to-mouth households.

I share his measurement strategy: I use the SCE density forecast to measure perceived risk and compare it with the realised process in the SIPP same-job stayer panel. I also corroborate his central finding: perceived risk lies below the full-information benchmark in every income group, and Section 3 interprets this gap as an upper bound. I depart from his analysis in four respects. First, I organise the gap by income using the matched same-job design and show that, once the realised process is decomposed, the discrepancy is two-sided: the bottom group over-perceives transitory risk, while the middle and top groups under-perceive permanent risk. The income gradient therefore places different groups on opposite sides of their group-specific benchmarks, even though perceived total risk remains below the benchmark for all groups. Second, I document the within-person dynamics of perceived risk, including its persistence and its response to realised shocks, whereas his evidence is cross-sectional. Third, I estimate an explicit updating rule, while he remains agnostic about belief formation by design. Fourth, I propose a microfounded explanation based on how pay is delivered: I measure the labelled share of transitory pay from the itemised pay components in the same records, whereas he attributes the level gap to unobserved heterogeneity or overconfidence. The economic incidence differs accordingly. In his model, belief heterogeneity affects aggregate wealth and its dispersion; here, the wedge is concentrated at the top because the saving-relevant distortion is under-perceived permanent risk, whose gradient is largest at high incomes.

The paper also connects to several strands of the literature on subjective income beliefs. Rozsypal and Schlafmann (2023) document an income-graded bias in expected income levels, whereas I study the second moment. Ben-David et al. (2018) show using the same survey that uncertainty about own-income growth falls with income and education. Koşar and van der Klaauw (2025) characterise the dynamics of the SCE's earnings expectations; I add same-job conditioning, a matched realised process, and an estimated updating rule. Using linked Danish survey and register data, Caplin et al. (2026) likewise find that survey-reported earnings risk is several times below its indirectly estimated counterpart. They attribute most subjective earnings risk to job transitions and find relatively small and symmetric risk conditional on remaining in the same job. Their instrument elicits expectations separately conditional on each possible employment transition, whereas the SCE question studied here conditions on remaining in the same job. Their measure therefore captures the risk associated with employment transitions, while mine isolates the risk that remains within the job. Koşar and Melcangi (2025) estimate how the marginal propensity to consume responds to subjective earnings uncertainty, directly measuring the behavioural margin through which the belief wedge in Section 6 operates. The proposed mechanism also relates to the "insurance versus information" distinction in the partial-insurance literature (Blundell et al., 2008). Pistaferri (2001) and Kaufmann and Pistaferri (2009) use subjective expectations to distinguish anticipated from unanticipated shocks and show that superior information can explain part of the perceived-realised gap. The mechanism here provides a related labelling channel, but one that is observable and varies systematically with income.

The updating rule builds on Pischke (1995), who frames inference about permanent income as a signal-extraction problem. The under-reaction implied by my updating rule is closely related to the information rigidity of Coibion and Gorodnichenko (2015), which Carroll et al. (2020) incorporate into consumption dynamics as sticky expectations. It contrasts with the diagnostic over-reaction of Bordalo et al. (2018), which Sanchez and Song (2025) embed in an incomplete-markets model. These mechanisms place perceived risk on opposite sides of the full-information benchmark, and the shortfall documented in Section 3 is more consistent with under-reaction. Experience-based learning (Malmendier and Nagel, 2011) provides another source of heterogeneity in beliefs, but its key state variable is age, whereas the gradient documented here rises with income conditional on demographics.

Finally, the proposed mechanism draws on documented differences in how pay is delivered across the income distribution: performance pay is concentrated among high-income workers (Lemieux et al., 2009), high-income earnings shocks are predominantly transitory (Guvenen et al., 2021), while pay is more volatile and fragmented at the bottom (Hardy and Ziliak, 2014; Morduch and Schneider, 2017).

2 Data and Measurement

The paper compares what workers believe about their wage risk with the risk their wages actually carry. I measure perceived risk using the New York Fed's Survey of Consumer Expectations (SCE) and realised risk using the Survey of Income and Program Participation (SIPP). Both are monthly within-respondent panels, with the SIPP covering reference years 2013–2023 and the SCE covering 2013–2025. I impose the same basic conditioning in both datasets, requiring a worker to remain in the same job with the same employer for at least six consecutive months, and apply survey weights throughout. Respondents are divided into three groups by median annual household income: bottom, below \$50k; middle, \$50k–\$100k; and top, at or above \$50k.[1]

The perceived side comes from the SCE density question, which asks each respondent for the distribution of her own earnings growth over the next twelve months *"on the same job, same hours, same employer"*. Responses are reported as a ten-bin histogram, following the density-forecast tradition of Engelberg et al. (2009). This conditioning makes the response useful for isolating wage risk: holding the job, hours, and employer fixed removes predictable changes associated with the worker's wage path, as well as variation in earnings arising from changes in labour supply. The elicited object is therefore closer to the conditional variance of the stochastic component of wage growth than to total income risk. It also excludes the risk of losing or changing jobs, which constitutes a substantial component of income risk in Caplin et al. (2026); the model in Section 6 reintroduces this margin separately through its employment block. From each reported histogram, I calculate the bin-midpoint variance as the perceived variance $\bar{\sigma}^2_{\text{perc}}$, exclude respondents who place all mass in a single bin, trim at 1%/99%, and drop each respondent's reports from the first wave in which she records a job change.

The realised side comes from the SIPP, which records earnings monthly and provides a unique employer identifier, allowing me to impose the same-job and same-employer conditions at the frequency targeted by the SCE question. Monthly earnings comprise regular wage and salary income together with bonuses, commissions, tips, and overtime, recorded as the amounts actually

[1]The cut-offs are nominal and fixed across the sample. The SCE skews towards higher incomes, so group sizes are uneven (estimation sample: 1,531 / 7,501 / 11,331 observations). The \$50k cut-off is chosen because finer income brackets become too thin to support the estimation (Appendix E.1), although a finer partition could locate more precisely where the gradient emerges.

paid; I exclude the self-employed. Unlike the SCE measure, these earnings also vary with hours actually worked. In Appendix C.7, I isolate the contribution of this difference in conditioning and find that accounting for it strengthens rather than overturns the main findings. I drop January observations, where reported cross-wave changes bunch spuriously, and trim squared changes at 1%/99%.

I residualise both objects on standard demographic controls before computing any moments, following Wang (2023): age and its square, sex, education, state, and year-month fixed effects, with the latter absorbing aggregate shocks common across groups. On the perceived side, this yields a residualised measure of perceived variance, $\bar{\sigma}^2_{\text{perc}}$; on the realised side, it yields a residualised log wage. The residualisation removes systematic variation predictable from demographics, location, and aggregate time effects, leaving the stochastic component on which the comparison is based. Finally, the SCE specification includes a survey-tenure term to absorb the panel conditioning associated with repeated subjective reports (Kim and Binder, 2023); no corresponding term is included on the realised-pay side.

3 Perceived and Realised Risk

This section sets out the facts the rest of the paper must explain. Perceived risk is steeply graded by income, and so is its persistence within person. Neither gradient has a counterpart in the realised wage process, which I estimate on matched same-job spells in the second half of the section. Realised wage shocks do, however, move beliefs over time.

3.1 Levels

Perceived variance is steeply graded by income, with the bottom of the distribution far above a middle and top that sit together. As Figure 1 shows, the weighted group means are 20.6 at the bottom, 11.3 in the middle, and 11.4 at the top. A worker at the bottom thus perceives about 1.8 times the wage uncertainty of a worker at the top. Ben-David et al. (2018) document in the same survey that uncertainty about own income growth falls with income and education. Here, the matched-stayer conditioning and the realised counterpart below sharpen that pattern into a statement about wage risk.

The gradient takes the form of a step: the middle and top groups are statistically indistinguishable under the midpoint measure and track each other closely, while the gap between the bottom and the other two groups survives changes in the income cut-off. The gradient also survives conditioning on numeracy and inflation uncertainty, both of which are themselves steeply graded by income (Appendix C.4). The bottom series drifts upward over the sample, which may reflect the fixed nominal cut-offs or the greater cyclical exposure of low-income pay.

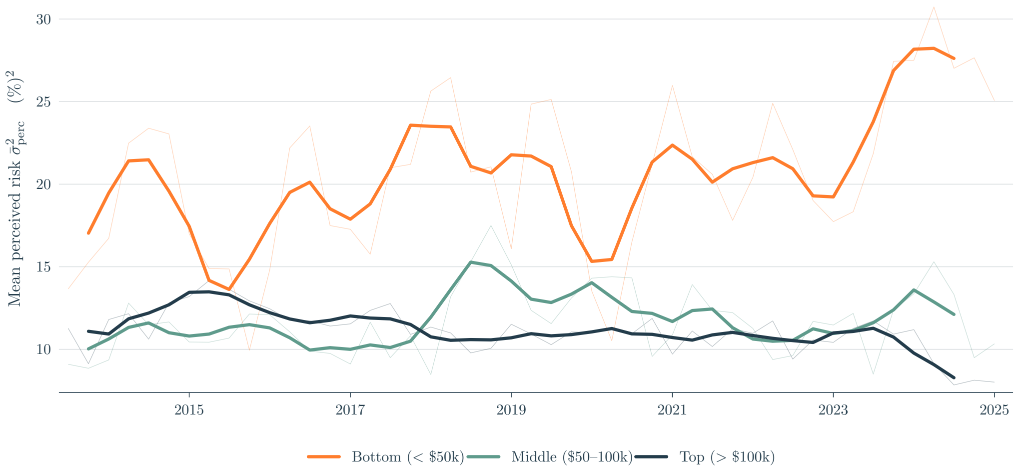

Figure 1: Group-mean perceived variance $\bar{\sigma}^2_{\text{perc}}$ by income group: quarterly means (thin) and their 4-quarter centred moving average (thick). Matched-stayer, trimmed and residualised SCE panel as in Section 2.

3.2 Persistence

Perceived risk is also more persistent within person at the bottom than at the top. I estimate a within-respondent AR(1) in perceived variance $\bar{\sigma}^2_{\text{perc}}$, with persistence coefficient $\hat{\phi}_g$, using Blundell–Bond system GMM (Blundell and Bond, 1998). The short, wide SCE panel makes the within estimator severely Nickell-biased, so I instrument the lagged dependent variable with deeper lags to address this bias, following Ma et al. (2020), who uses the same estimator for a short-panel setting.

	Monthly		Quarterly	
Group	$\hat{\phi}_g$	SE	$\hat{\phi}_g$	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 $\hat{\phi}_g$ of perceived variance, Blundell-Bond system GMM. Instruments: $\bar{\sigma}^2_{\text{perc}}$ 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.

The bottom group is the most persistent of the three, and equality with the top group is rejected at both monthly and quarterly cadences. This gradient does not appear to be inherited from the realised process: in Appendix C.5, I show that realised volatility instead becomes more persistent with income. Nor is the gradient an artefact of the reporting cadence. Classical measurement noise in the monthly reports biases monthly persistence estimates downward and is partly averaged out at the quarterly frequency, pushing the quarterly estimates upward. Noise alone therefore cannot account for the gap, since doing so would require an implausibly persistent underlying process. At either cadence, the bottom group holds its beliefs for longer than the top group.

Taken together with the level differences, these persistence estimates imply that calibrating perceived income risk to realised wage data requires reproducing not only the income gradient in the level of perceived risk, but also the gradient in how persistently workers hold those beliefs.

3.3 The Realised Wage Process

To compare realised and perceived risk at the survey's horizon, I estimate the standard permanent-transitory decomposition of Meghir and Pistaferri (2004) on matched same-job spells:

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

Here $w_{i,t}$ is the residualised log wage, $z_{i,t}$ the predictable Mincerian component, 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 $\sigma^2_{\psi,g}$, where g indexes the worker's income group, and accumulates through the random walk $p_{i,t}$. Economically, it captures shocks such as merit raises or cost-of-living adjustments. The transitory shock $\theta_{i,t}$ is iid with variance $\sigma^2_{\theta,g}$ and captures temporary fluctuations such as discretionary bonuses or unpredictable working hours.

The SCE question targets this stochastic component. Because it conditions on *the same job, hours, and employer*, the predictable component is held fixed, so the elicited variance corresponds to

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

On the SIPP side, I residualise wages on demographics to remove z and then difference the residual to recover Δe directly. The recorded wage is monthly earnings per average weekly hour reported for the month, so it includes tips, overtime, and premium pay. Both measures therefore target the stochastic component of wage changes within the same employer, with the qualification that the SCE also conditions on hours.

For the baseline realised moments, I retain all months, whether or not reported hours change. A persistent change in scheduled hours is itself a source of income risk borne by the worker, even though the SCE question conditions it away. I examine the alternative conditioning on unchanged hours separately in Appendix C.7.

Two moments per income group identify the two wage primitives. The one-period change is

$$\Delta e_t = \psi_t + \theta_t - \theta_{t-1}, \tag{3}$$

and let $\bar{S}_{m,g} \equiv \mathbb{E}[(\Delta e)^2]$ denote its variance. Because ψ and θ are independent and iid,

$$\bar{S}_{m,g} = \sigma^2_{\psi,g} + 2\sigma^2_{\theta,g}, \qquad \mathrm{Cov}(\Delta e_t, \Delta e_{t-1}) = -\sigma^2_{\theta,g}. \tag{4}$$

The first-order autocovariance therefore identifies the transitory variance, while the remaining variance identifies the permanent variance, separately for each income group. Autocovariances beyond the first lag are close to the zero restrictions implied by the model, and conditioning on stayers does not materially change the estimates (Appendix C.6).

Group	$\sigma^2_{\theta,g}$ (transitory)	$\sigma^2_{\psi,g}$ (permanent)	$\bar{S}_{m,g}$
Bot	2.52	11.89	16.92
Mid	4.42	11.60	20.44
Top	5.83	11.22	22.88

Table 2: SIPP permanent-transitory identification by income group, monthly cadence, in $(\%)^2$. Moments are computed from unrounded inputs and rounded for display.

The realised gradient runs, if anything, in the opposite direction to the perceived gradient (Table 2). Permanent variance is approximately flat across income groups, while transitory variance rises with income, plausibly reflecting the greater importance of bonus pay among higher-income workers. Conditioning on unchanged reported hours sharpens rather than overturns this pattern: both components then rise mildly with income (Appendix C.7). The contrast with the perceived gradient therefore does not depend on the treatment of hours.

3.4 The FIRE Benchmark

The perceived and realised objects meet at the twelve-month horizon targeted by the survey. Rolling the monthly process forward twelve months under iid shocks (derivation in Appendix A) gives

$$\mathrm{Var}(e_{t+12} - e_t) \;=\; 12\sigma^2_{\psi,g} + 2\sigma^2_{\theta,g}, \tag{5}$$

which is the twelve-month risk a full-information rational-expectations (FIRE) agent would report. This benchmark is around $148\,(\%)^2$ for every group, and falls to 79–98 under the same-hours conditioning in Appendix C.7. Perceived variance is substantially below either benchmark, at 20.6, 11.3, and 11.4, corresponding to between one-quarter and one-thirteenth of the benchmark (Table 4, Appendix A). The gap between perceived and realised income risk is similar in magnitude to that documented by Wang (2023).

How much of this gap reflects genuine miscalibration is difficult to say, because the benchmark itself is fragile. Equation (5) multiplies the permanent variance by twelve, while SIPP earnings are self-reported: the classical component of reporting error is treated as transitory, but any persistent component is indistinguishable from true permanent shocks and therefore enters the benchmark twelvefold (Bound and Krueger, 1991). Administrative payroll data, which are free from reporting error, place job-stayers' base-wage dispersion well below the benchmark (Grigsby et al., 2021). Reporting error is plausibly common across groups, however, so its effect should largely cancel in the cross-group comparison. I therefore interpret the level gap only as an upper bound and base the argument primarily on the contrast across groups.

3.5 Belief Dynamics and Realised Shocks

The cross-section rules out the realised process as an explanation for the income gradient, but realised shocks can still shape belief dynamics. Figure 2 shows that pooled perceived risk co-moves with the absolute deviation of aggregate wage growth from its sample mean. The pattern is consistent with wage-growth surprises raising perceived risk over time.[2]

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

Together, these facts impose two requirements on any belief rule. Perceived variance must be allowed to move over time, since the dynamics rule out constant beliefs, and realised shocks must be able to feed into those movements. Whether the response to a largely common signal differs systematically across income groups is an empirical question, which I take up using within-person variation in Section 5.

4 Belief Formation under Illegible Pay

This section develops a model of how workers infer income risk from the pay they observe, and what this implies for saving. The worker sees one pay change each period, which combines permanent and transitory shocks, and must infer how much of the variation reflects a lasting change in earning power. Some transitory pay arrives labelled, for example as a bonus or commission on a separate payslip line, and the worker can *net* it out before forming her belief. The remainder mixes permanent and transitory variation, and she cannot decompose it any further.

The mechanism generates four results. First, labelling makes the signal about permanent earning power cleaner, but the worker does not fully adjust her attribution rule to account for this improvement. She therefore assigns too little of each observed surprise to the permanent component, and her perceived permanent risk falls below the true variance. Second, the same netting reduces the transitory variation left in the signal, pushing perceived transitory risk downward. Third, noise in reading the unlabelled remainder is itself indistinguishable from a transitory shock and therefore raises perceived transitory risk. Total perceived risk can consequently lie either above or below realised risk: the permanent component is always under-perceived, while the transitory component can be either over- or under-perceived. Fourth, the two errors have different behavioural

[2] I use absolute deviations because a realised shock of either sign should increase perceived uncertainty; squared deviations produce a very similar pattern. The comovement motivates the shock regressor used in Section 5, where the within-person estimates provide the evidence on whether realised shocks move beliefs.

consequences. Permanent-risk misperception affects precautionary saving, whereas transitory-risk misperception is largely buffered within the period. The saving distortion therefore need not be largest for the worker with the largest total belief error.

Two primitives govern the mechanism. The first is pay legibility, ρ_g, the share of transitory pay that arrives sufficiently identified for the worker to net it out before updating. The second is reading noise, ω_g^2, the variance of the error with which she interprets the remaining pay. I treat both as primitives of the belief process and measure them in Section 5. The labelled share is constructed from the itemised components of the same pay records used to estimate the realised wage process, while reading noise is identified from the component of the transitory prior left unexplained by the measured share.

4.1 Environment

A worker's wage follows the permanent-transitory process of Section 3, $\Delta e_t = \psi_t + \theta_t - \theta_{t-1}$, where ψ_t is an iid permanent shock with variance $\sigma^2_{\psi,g}$ that accumulates in a random walk, and θ_t is an iid transitory shock with variance $\sigma^2_{\theta,g}$ that disappears within the month. The worker knows the form of this process but must learn the variance of permanent earning power from the pay she observes.

The worker observes one pay change per period. Since permanent and transitory shocks enter the same pay change, the raw squared change contains both components: $\mathbb{E}[(\Delta e)^2] = \sigma^2_{\psi,g} + 2\sigma^2_{\theta,g} \equiv \bar{S}_{m,g}$. Her problem is therefore a signal-extraction problem in the spirit of Pischke (1995). She is not trying to infer the level of permanent income from noisy observations, but rather the variance of the permanent component from observations that also contain transitory variation. A household may therefore track the level of its persistent income well, as Druedahl and Jørgensen (2020) infer from consumption responses, while still misperceiving the variance of its income shocks.

Pay is not necessarily observed as one undifferentiated number. Let $\rho_g \in [0, 1]$ denote the share of transitory pay that arrives labelled as transitory before the worker updates. A labelled payment might appear as a separate line on a payslip, be announced in advance with its amount, or recur at a predictable time of year. The worker can therefore net the labelled component out of the pay change and update only on what remains. Labelling is plausibly more common higher in the income distribution, where performance pay is often explicitly identified (Lemieux et al., 2009). It can also arise through other institutional arrangements. The measure used in Section 5 captures only labelling that is visible in the itemised pay records, so it should be interpreted as a lower bound on the information available to workers.

For intuition, consider a salaried worker whose pay is a few hundred euros higher in one month because of a commission. If the commission appears as a separate line on the payslip, she can identify the source of the increase and net it out. The remaining salary change is close to zero, so the observation contains little news about her permanent earning power. By contrast, when the same increase arrives through a combination of variable hours, tips, and irregular supplements, the worker may not be able to identify which part of the change is transitory.

Write the transitory shock as $\theta_t = \theta_t^l + \theta_t^u$, where the labelled and unlabelled components have variances $\mathrm{Var}(\theta_t^l) = \rho_g \sigma^2_{\theta,g}$ and $\mathrm{Var}(\theta_t^u) = (1 - \rho_g)\sigma^2_{\theta,g}$. After netting out the labelled component, the residual signal has variance

$$V_{u,g} = \sigma^2_{\psi,g} + 2(1 - \rho_g)\sigma^2_{\theta,g} \leq \bar{S}_{m,g}, \tag{6}$$

with equality when no transitory pay is labelled. Greater legibility therefore reduces the variance of the signal the worker updates on while leaving its permanent component unchanged. The residual is therefore a cleaner signal of permanent earning power.

The second primitive is the worker's ability to read the unlabelled remainder. Consider a worker whose pay varies with shift counts, cash tips, and irregular supplements, potentially alongside changes in hours (Hardy and Ziliak, 2014; Morduch and Schneider, 2017). She may observe that monthly pay has increased without being able to determine precisely how much of the increase reflects her wage rather than other components of earnings. In Appendix D.5 I show that within-person hours volatility indeed falls monotonically with income.

I capture this uncertainty as reading noise:

$$\tilde{y}_t = p_t + \theta^u_t + \nu_t, \qquad \nu_t \overset{\text{iid}}{\sim} (0, \omega^2_g), \tag{7}$$

where p_t is the permanent component and ν_t is the worker's reading error. The error enters the observed change as $\Delta\nu_t$. It therefore has variance $\text{Var}(\Delta\nu_t) = 2\omega^2_g$, and first-order autocovariance $\text{Cov}(\Delta\nu_t, \Delta\nu_{t-1}) = -\omega^2_g$, with zero autocovariance at higher lags. Its statistical signature is therefore identical to that of a transitory shock. The worker cannot distinguish the two from her own observations, so reading noise is recorded as transitory risk in her beliefs. I am not taking a stand on specifying the source of this error since misreading pay components, uncertainty about hours, and differences in reporting or interpretation all generate the same object in the worker's information set.

4.2 Modelling Choices

The survey provides one perceived-risk series per worker, while the model contains two latent beliefs, one about permanent risk and one about transitory risk. The two cannot therefore be separately estimated from the perceived-variance series alone. I make two modelling choices to assign the dynamics and levels to these components.

First, I let the permanent-risk belief evolve over time, $\hat{\sigma}^2_{\psi,t}$, while holding the perceived transitory variance at a group-specific prior, $\sigma^2_{\theta,\text{perc},g}$. I made this decision since permanent shocks accumulate in the worker's earning power and therefore alter the entire future path of consumption, whereas transitory shocks disappear within the period and are largely absorbed by the household's buffer. The belief with the stronger saving implications should therefore be the one that responds to new information. A moving belief about transitory risk would add another latent state without much behavioural content in the buffer-stock economy. Moreover, the reading error in (7) is observationally equivalent to a transitory shock, which makes the transitory prior the natural place for uncertainty that the worker cannot resolve. This restriction comes at a cost since the perceived-variance series alone cannot test whether the transitory belief is actually static. Any belief-side misallocation between the permanent and transitory components is therefore absorbed by the attribution structure of the model.

Second, I hold fixed the share of each squared surprise that the worker attributes to permanent earning power. Let λ_g denote this attribution share. The worker nets out labelled transitory pay before updating, so the remaining signal is cleaner than the raw pay change. A Bayesian worker would respond to this improvement in signal quality by increasing the share of the remaining surprise attributed to permanent risk. I instead assume that the worker applies the same attribution rule before and after netting.

This distinction is central to the mechanism. If the worker fully adjusted her attribution to the cleaner signal, labelling would improve information without creating a persistent bias in the level of perceived permanent risk. Under the constant-attribution rule, by contrast, the worker continues to treat the cleaner residual as though it contained the same mixture of permanent and transitory variation as the original pay change. Netting therefore reduces the amount of variation she attributes to permanent risk, generating a permanent-risk belief below the realised variance.

The constant attribution rule has a precedent in adaptive learning. For a random-walk state, a constant-gain updating rule is the standard perpetual-learning benchmark (Muth, 1960) and underlies much of the adaptive-learning literature (Evans and Honkapohja, 2001; Malmendier and Nagel, 2016). The restriction here is not the use of a constant gain itself, but the failure to adjust the attribution share when the information content of the signal changes. This is consistent with evidence that agents compress the informativeness of signals, updating too little on highly informative signals and too much on weak ones (Edwards, 1968; Augenblick et al., 2025). Noisy cognition can generate similar compression (Enke and Graeber, 2023).

4.3 The Belief Rule

The permanent-risk belief follows an exponentially weighted moving average of past squared wage changes:

$$\hat{\sigma}^2_{\psi,t} = \underbrace{(1-\mu)}_{\text{persistence}} \; \hat{\sigma}^2_{\psi,t-1} + \underbrace{\mu}_{\text{update weight}} \; \underbrace{\lambda}_{\text{attribution share}} \; (\Delta e_{t-1})^2. \tag{8}$$

where $\mu_g \in [0,1]$ is the update speed and $\lambda_g \in [0,1]$ is the share of each squared surprise that the worker attributes to permanent earning power. A higher μ_g makes the belief respond more strongly to recent news, while a higher λ_g raises the amount of each surprise attributed to the permanent component.

The recursion is a constant-gain learning rule. Its persistence is $1-\mu_g$, so a shock to perceived risk decays geometrically over time. In squared shocks, it is analogous to an integrated GARCH(1,1) process. It can also be interpreted as a constant-gain approximation to a Kalman filter for learning about an evolving variance. Appendix D.1 derives the corresponding signal-extraction problem and relates the gain to the signal-to-noise ratio of the squared wage changes.

The survey asks about twelve-month earnings growth, so the worker maps her beliefs about the monthly permanent and transitory components into a twelve-month perceived variance:

$$\bar{\sigma}^2_{\text{perc},t} = 12\,\hat{\sigma}^2_{\psi,t} + 2\sigma^2_{\theta,\text{perc}}, \tag{9}$$

The permanent component accumulates over twelve independent monthly shocks, while the transitory component enters at the two endpoints of the twelve-month change. This is the same aggregation used to construct the full-information benchmark in (5).

The learning rule also has a useful stationary interpretation. Taking unconditional expectations of (8) gives $\mathbb{E}[\hat{\sigma}^2_\psi] = \lambda\, \bar{S}_{m,g}$ so the permanent-risk belief is unbiased on average when $\lambda^*_g = \sigma^2_{\psi,g}/\bar{S}_{m,g}$ Using the realised moments in Table 2, this full-information attribution share is 0.70, 0.57, and 0.49 for the bottom, middle, and top groups, respectively.

4.4 Legibility and Beliefs

I now connect how pay is delivered to the beliefs generated by the learning rule. The first result is that labelling reduces the share of an observed pay change that the worker attributes to permanent earning power. Labelled transitory pay contributes to the observed change but carries no information about permanent earning power. When the worker nets it out, the remaining change is smaller, but under the constant attribution rule she does not increase the share she assigns to permanent risk.

Proposition 1 (Attribution under partial labelling)**.** *Suppose ψ and θ are jointly Gaussian and the labelled and unlabelled pieces are independent. Then the share of the full squared change a worker books as permanent is*

$$\lambda_g \;=\; \frac{\sigma^2_{\psi,g}\big(\sigma^2_{\psi,g} + 2(1-\rho_g)\sigma^2_{\theta,g}\big)}{(\sigma^2_{\psi,g} + 2\sigma^2_{\theta,g})^2} \;=\; \lambda^*_g\,\frac{V_{u,g}}{\bar{S}_{m,g}}, \tag{10}$$

which is strictly decreasing in ρ_g and equals the full-information share λ^*_g at $\rho_g = 0$.

The proof is in Appendix B. The labelled component is part of the observed pay change but is removed before the worker forms her permanent-risk signal. As ρ_g rises, a larger part of the observed change is therefore uninformative about permanent earning power. Measured against the full wage change, the share attributed to permanent risk falls. This delivers a simple prediction: if pay becomes more legible as income rises, the attribution share should fall with income.

Proposition 2 (The permanent-belief discount)**.** *Under the constant attribution maintained above, the belief* (8) *rests at*

$$\mathbb{E}[\hat{\sigma}^2_\psi] \;=\; \lambda^*_g\, V_{u,g} \;=\; \sigma^2_{\psi,g}\,\frac{V_{u,g}}{\bar{S}_{m,g}} \;\le\; \sigma^2_{\psi,g}, \tag{11}$$

with equality only at $\rho_g = 0$, and the discount $V_{u,g}/\bar{S}_{m,g}$ is strictly decreasing in ρ_g. The resting belief does not depend on the reading noise ω^2_g.

The resting permanent-risk belief is therefore the realised permanent variance multiplied by the fraction of pay variance that remains after labelled transitory pay is removed. More legible pay produces a cleaner signal, but under the constant attribution rule the worker does not adjust her attribution share to account for this improvement. The resulting belief is consequently too low.

The absence of ω^2_g from the permanent belief follows from a cancellation. Averaging the recursion over a stationary state leaves the resting belief at the attribution times the mean of the signal. Reading noise inflates the mean of the signal, while the attribution is rescaled to that mean, so the noise cancels from the stationary permanent belief and survives only in the transitory prior.[3]

The second result concerns the transitory component. Netting labelled pay lowers the amount of transitory risk that remains in the worker's signal, while difficulty in reading the unlabelled remainder adds noise that is observationally equivalent to a transitory shock. The two forces therefore operate in opposite directions.

4.5 The Perceived-Risk Wedge

The previous results determine how legibility affects the two components of perceived risk. I now combine them to characterise the gap between perceived and realised twelve-month risk.

Define the perceived-risk wedge as the difference between the agent's perceived and realised second moments, using the weights from the twelve-month aggregation:

$$\Delta\sigma^2 \;=\; 12\big(\hat{\sigma}^2_{\psi} - \sigma^2_{\psi,g}\big) \;+\; 2\big(\sigma^2_{\theta,\text{perc},g} - \sigma^2_{\theta,g}\big). \tag{12}$$

Proposition 3 (Two-sided perceived risk)**.** *The worker's transitory prior is her netting floor plus her reading noise,*

$$\sigma^2_{\theta,\text{perc},g} \;=\; (1-\rho_g)\,\sigma^2_{\theta,g} \;+\; \omega^2_g, \qquad \omega^2_g \geq 0, \tag{13}$$

so the permanent gap $\hat{\sigma}^2_{\psi} - \sigma^2_{\psi,g} \;=\; \sigma^2_{\psi,g}\big[V_{u,g}/\bar{S}_{m,g} - 1\big] \;\leq\; 0$ *is one-signed, while the transitory gap* $\sigma^2_{\theta,\text{perc},g} - \sigma^2_{\theta,g} = \omega^2_g - \rho_g\sigma^2_{\theta,g}$ *takes either sign. Total perceived risk can therefore sit on either side of the truth.*

The permanent component is always under-perceived because labelling removes transitory variation from the signal without inducing the worker to increase the attribution share. The transitory component is different: labelling reduces the transitory variation that remains in the worker's signal, while reading noise adds variation that she cannot distinguish from a transitory shock. Hence, labelling pushes perceived transitory risk down, whereas reading noise pushes it up.

This two-sidedness is the key implication of the mechanism. A single directional source of belief error cannot generate a cross-section in which some groups over-perceive risk while others under-perceive it. Here the two forces operate in opposite directions: where pay is fragmented and difficult to read, reading noise can dominate and push perceived transitory risk above its realised counterpart; where pay is more legible, netting can dominate and push it below.

The crossing therefore occurs through the transitory component. The permanent component is always under-perceived, while the transitory component can lie on either side depending on the relative strength of legibility and reading noise. This allows total perceived risk to lie either above or below realised risk.

[3] The implemented reading is the third of three that share this resting mean. The agent updates on the squared change of her own noisy netted reading (7), whose mean is $V_{u,g} + 2\omega^2_g = \sigma^2_{\psi,g} + V_x$ with $V_x = 2\sigma^2_{\theta,\text{perc},g}$, and the attribution is the rescaled $\tilde{\lambda}_g = \sigma^2_{\psi,g} V_{u,g}/[\bar{S}_{m,g}(\sigma^2_{\psi,g} + V_x)]$, chosen so that the resting mean is unchanged (Appendix D.1, Table 11): the reading error inflates her signal and the rescaling offsets it, so it never reaches the resting permanent belief.

4.6 Prudence and Incidence

The perceived-risk wedge matters for behaviour through the curvature of the value function. Because the perceived and realised shocks have the same means, a second-order approximation isolates the effect of the difference in their variances:

$$\widetilde{\mathbb{E}}\,V - \mathbb{E}\,V \;\approx\; \frac{1}{2}\big(\hat{\sigma}^2_{\psi} - \sigma^2_{\psi,g}\big)V_{\psi\psi} + \frac{1}{2}\big(\sigma^2_{\theta,\text{perc},g} - \sigma^2_{\theta,g}\big)V_{\theta\theta}, \tag{14}$$

where $\widetilde{\mathbb{E}}$ and $\mathbb{E}$ denote expectations under the worker's perceived and the true distributions, respectively, and $V_{\psi\psi}$ and $V_{\theta\theta}$ are the corresponding curvatures evaluated at the mean. The equation identifies the channel through which a belief about risk enters the household problem, while the quantitative magnitude is to be determined by the model in Section 6.

The same channel can be seen from the Euler equation. A belief about future risk enters through the expected marginal utility of next-period consumption, $\widetilde{\mathbb{E}}[u'(c')]$. Expanding around mean next-period consumption gives

$$\widetilde{\mathbb{E}}[u'(c')] - \mathbb{E}[u'(c')] \;\approx\; \frac{1}{2}g_p\mathbb{E}\Big[u'''(c')c_y^2 + u''(c')c_{yy}\Big], \qquad g_p \equiv \hat{\sigma}^2_{\psi} - \sigma^2_{\psi,g}, \tag{15}$$

where c_y and c_{yy} are the first two derivatives of the consumption policy with respect to income. Under the CRRA preferences used below, the leading term is positive, so an agent who perceives less permanent risk than she faces under-saves for precaution.

The distinction between permanent and transitory risk is important for the incidence of the wedge. A transitory shock is largely absorbed within the period by the household's buffer, whereas a permanent shock changes the consumption path going forward. The variance error therefore matters much more when it concerns permanent rather than transitory risk. This is particularly important when the behavioural margin is precautionary saving: the permanent-risk wedge affects the desired buffer, while the transitory wedge has a comparatively small effect.

Proposition 4 (Incidence)**.** *Let g_p denote the saving-relevant gap in* (15)*, and let $a'(m, \hat{\sigma}^2_{\psi})$ denote the saving policy. Then*

1. *$\partial g_p/\partial \rho_g < 0$ and $\partial g_p/\partial \omega^2_g = 0$: the saving-relevant distortion depends on legibility but not on reading noise;*
2. *only the permanent gap materially affects saving: a transitory belief error is buffered within the period and enters the consumption response at second order;*
3. *$\partial a'/\partial \hat{\sigma}^2_{\psi} = 0$ at the borrowing constraint $a' = \underline{a}$ and is strictly positive only when the household has slack, whereas a level error reaches a constrained household one for one.*

The first part follows from the cancellation in Proposition 2. Reading noise is recorded as transitory risk and therefore does not enter the saving-relevant permanent gap. Legibility, by contrast, lowers the permanent-risk belief and hence increases the gap between perceived and realised permanent risk.

The second part follows from how the two shocks affect consumption. A household with slack can absorb a transitory draw using its buffer, so next-period consumption responds only weakly to it. A permanent shock instead shifts the entire future income path, making its effect on consumption much larger. Since a variance error enters through the squared consumption response, the transitory component has little saving relevance relative to the permanent component.

The third part concerns the role of the borrowing constraint. A constrained household already saves at its lower bound, so changing perceived risk does not change its actual saving. A level error is different: it changes the income that arrives and therefore reaches consumption one for one even when the constraint binds. If the constraint binds more often at low incomes, it further mutes the behavioural response of the bottom group.

Corollary 1 (The gradient in behaviour runs opposite to the gradient in beliefs)**.** *Wherever legibility rises with income, the behavioural wedge rises with income. It is therefore largest for the group whose reported risk is lowest and smallest for the group whose reported risk is highest and most distorted.*

The corollary follows from the composition of the belief errors. The bottom group has the largest total distortion, but much of it comes from an inflated transitory prior that has little saving relevance. The top group reports less total risk, but its error is concentrated in the permanent component, which directly affects precautionary saving. The result is therefore that the gradient in behavioural distortions runs opposite to the gradient in reported risk.

To illustrate the incidence, consider a low-income household that is close to the borrowing constraint. Its perceived-risk error may be large, but if that error is primarily an overstatement of transitory risk, it has little effect on saving; and even an error in permanent risk has little effect while the constraint binds. A higher-income household may report less total risk, but if it has slack and under-perceives permanent risk, it responds by reducing its precautionary buffer. The largest belief distortion therefore need not generate the largest behavioural distortion: incidence depends both on which component is misperceived and on whether the household has room to adjust its saving.

4.7 Level Shocks versus Risk Shocks

The belief rule responds to realised wage shocks, but a level shock and a risk shock move perceived risk differently. A level shock changes the realised wage in a given period and so raises the squared surprise the worker updates on. A risk shock instead changes the variance of future wage growth, and so changes the distribution of every surprise that follows.

Under the belief updating rule (8), a level shock raises perceived permanent risk on impact by $\mu_g \lambda_g$ per unit of the squared change, and the effect decays geometrically, to $\mu_g \lambda_g (1 - \mu_g)^h$ after h months. A one-off wage shock therefore lifts perceived risk only temporarily.

A persistent rise in $\sigma^2_{\psi,g}$ works differently. Squared wage changes stay larger on average, so the worker keeps revising upward. Under the constant attribution she books only a fraction of the rise, so a true increase of $\Delta\sigma^2_\psi$ raises her perceived permanent variance by $\lambda_g \Delta\sigma^2_\psi$ alone, short of the truth.

This distinction matters for the dynamics of precautionary saving. A recession that produces a temporary wage decline can raise perceived risk through the unusually large realised shock, but the effect fades as the worker receives subsequent observations. A persistent increase in income risk instead leaves households facing a sustained wedge between perceived and realised risk. The model therefore predicts a temporary belief response to level shocks but a persistent under-response to changes in the underlying risk process.

The mechanism can also vary over the business cycle. If downturns make pay less legible, for example because bonuses and commissions disappear while hours become more variable, the labelled share ρ_g falls and reading noise ω_g^2 rises. The resulting increase in perceived risk then reflects not only the lower level of income but also a change in how that income is delivered. A recession can therefore affect the perceived-risk wedge through two distinct channels: the realised wage shock changes beliefs directly, while changes in pay legibility alter the mapping from realised shocks into beliefs.

5 Estimation and Measurement

This section takes the belief rule to the data in two steps. I first estimate the learning rule within person by income group. I then measure the labelled share from itemised pay components and use it to test whether the mechanism can account for the estimated belief gradients.

The belief recursion (8) is written in the unobserved $\hat{\sigma}_{\psi}^2$. Combining it with the level identity (9) eliminates the latent belief and gives an AR(1) in the observable perceived variance:

$$\bar{\sigma}^2_{\text{perc},i,t} = \underbrace{2\mu_g \sigma^2_{\theta,\text{perc},g}}_{\beta_{\text{cons}}} + \underbrace{(1-\mu_g)}_{\beta_{\text{lag}}} \bar{\sigma}^2_{\text{perc},i,t-1} + \underbrace{12\mu_g \lambda_g}_{\beta_{\text{shock}}} (\Delta e_{i,t-1})^2 + \alpha_i + u_{i,t}, \tag{16}$$

where α_i is a person fixed effect. The three coefficients map one-to-one to the structural parameters:

$$\mu_g = 1 - \hat{\beta}_{\text{lag}}, \qquad \sigma^2_{\theta,\text{perc},g} = \frac{\hat{\beta}_{\text{cons}}}{2\mu_g}, \qquad \lambda_g = \frac{\hat{\beta}_{\text{shock}}}{12\mu_g}. \tag{17}$$

The person effects are normalised to have mean zero within each income group. Thus β_{lag} identifies the updating speed from within-person persistence, β_{cons} identifies the transitory prior from the group level, and β_{shock} identifies the response to realised wage news.

5.1 The Shock Regressor

The learning rule calls for the worker's own $(\Delta e_{i,t-1})^2$, which the SCE does not observe. I therefore use the SIPP cell mean of squared wage changes among workers in the same cell c = (income group × Census division × month). This captures wage news common to workers facing the same local labour-market conditions. Following Guerreiro (2023), I smooth the cell series with a three-month trailing moving average.

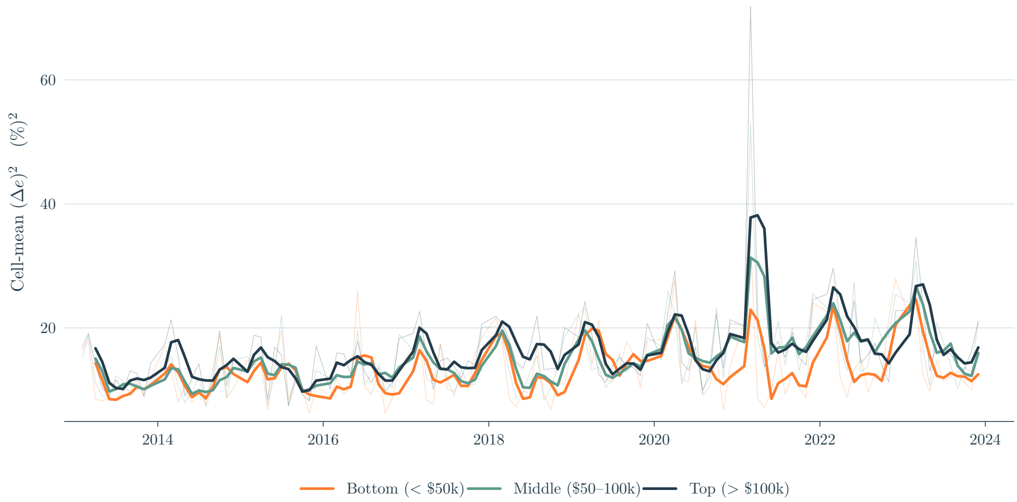

Figure 3: Cell-mean $(\Delta e_{i,t})^2$ by income group, national monthly, in $(\%)^2$. Thin lines: raw monthly cell means. Thick lines: the 3-month trailing moving average used as the regressor in (16).

Figure 3 shows that realised wage shocks are largest for the middle and top groups, consistent with their higher realised transitory variance and the concentration of discretionary pay at the top (Guvenen et al., 2021; Halvorsen et al., 2024). The bottom therefore does not receive larger shocks that could mechanically explain its stronger belief response. Its stronger response reflects how it updates to the shocks it receives.

The cell mean introduces measurement error because it captures only the component of a worker's wage shock that is common within her cell. This attenuates β_{shock}, and therefore $\hat{\lambda}$, towards zero. The coefficient is consequently best interpreted as the response of perceived risk to common local wage news. Appendix D.3 estimates the group-specific loadings at 0.43, 0.56, and 0.58 for the bottom, middle, and top, respectively. Correcting for these loadings widens the cross-group gradient, so the estimated gradient is a lower bound on the underlying gradient.[4]

5.2 Estimates by Income Group

Joint GMM estimation of (16) by income group gives the estimates in Figure 4, which I set against the realised SIPP transitory variance. I treat the shock as endogenous, since it shares survey timing and aggregate news with the belief, and instrument it with its own lags and lags of perceived variance, as in the Blundell–Bond estimation of Table 1. The full estimates are reported in Table 5 (Appendix C).

[4] The loadings measure the covariance between individual and cell-level wage shocks relative to the variance of the cell-level shock.

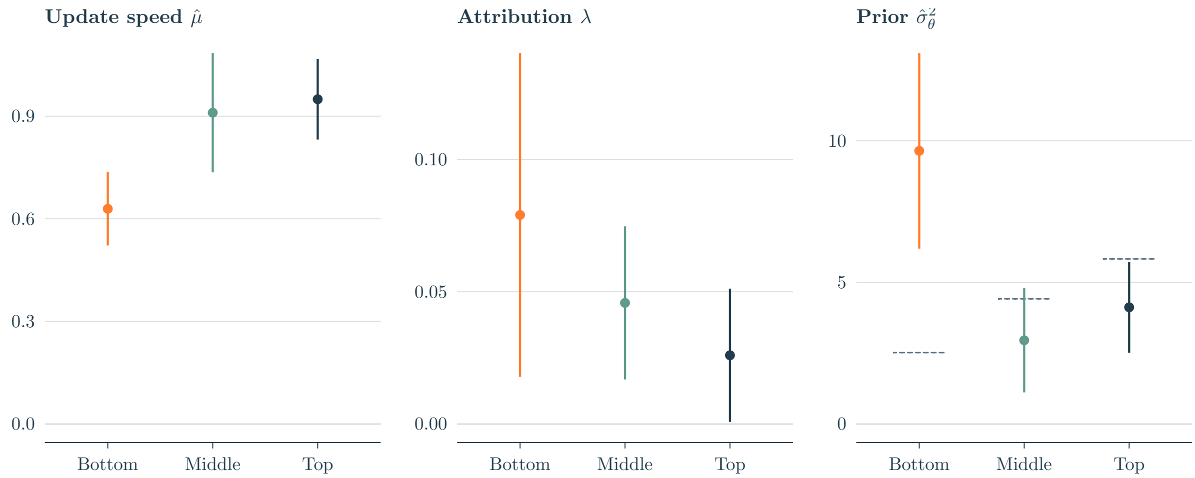

Figure 4: Headline estimates by income group: update speed μ_g, attribution λ_g, and transitory prior $\sigma^2_{\theta,\mathrm{perc},g}$, with 95% confidence intervals. Grey dashes in the prior panel mark realised transitory variance $\sigma^2_{\theta,g}$ (SIPP).

Update speed. The bottom updates significantly more slowly than the middle or top, matching the persistence ordering in Table 1. A given wage shock therefore remains in the bottom group's perceived risk for longer.

Attribution. All three $\hat{\lambda}_g$ estimates are positive and statistically significant, and the point estimates decline monotonically with income. The bottom's attribution share is about three times the top's. The cross-group differences are not statistically significant, however, and part of the decline is consistent with full information, since λ^*_g itself falls with income, although at roughly half the estimated rate.

Transitory prior. Realised transitory variance rises with income, while the perceived prior is almost four times its realised counterpart at the bottom and lies below it in the middle and at the top. The cross-group gradient is statistically significant. The high level of perceived risk at the bottom therefore comes primarily from its transitory prior, whose source is examined in the second half of this section.

5.3 Robustness

I first add an occupation × year-month fixed effect to the shock construction. The persistence and attribution gradients remain, with the latter becoming sharper, while the prior gradient retains its direction but loses cross-group significance (Appendix C.1). I do not use this specification as the baseline because occupation is itself part of the mechanism: low-income workers sort into occupations with structurally more uncertain pay (Morduch and Schneider, 2017), so controlling for occupation removes part of the variation through which pay legibility may operate.

The estimates are also stable to changes in the instrument lags, trimming, and COVID treatment (Appendix C.2). The persistence gradient survives all variants except the tighter within-group trim, where its significance becomes marginal, and the attribution gradient remains monotone throughout. The bottom's prior excess remains in every specification, ranging from 6.4 to 11.2

against a realised variance of 2.52, although its cross-group significance disappears under the COVID exclusion and the occupation control. I retain the COVID period because these years contain large realised wage movements that are part of the SIPP benchmark and provide identifying variation for the belief response.

Finally, I split the shock by sign, pooling income groups for power. Upside news receives roughly twice the weight of downside news. I treat this asymmetry as suggestive rather than as a central result, since pooling obscures which part of the income distribution drives it (Appendix C.3). If the asymmetry is actually present across all groups, I would interpret it as simple loss-aversion: workers appear to treat gains as more informative about persistent earning power than losses.

5.4 Measuring the Labelled Share

The three belief gradients are clear in the data, but their underlying mechanism remains to be established. The theory of Section 4 ties them to the labelled share ρ_g, which can be measured directly from the itemised pay components in the SIPP. I therefore use the same stayer panel that supplies the realised wage process to measure how much of each worker's pay movement arrives explicitly labelled as transitory.

The SIPP itemises the components of each paycheck, allowing me to measure the labelled share ρ_g directly on the same stayer panel used to estimate the realised income process. For each worker and month, I observe how much pay arrived as an identifiable component, such as a bonus or commission, and how much did not. I classify tips and overtime as unlabelled. The resulting share of labelled pay is 0.121 at the bottom, 0.178 in the middle, and 0.358 at the top. Appendix D.2 gives the construction and specification checks, which produce similar estimates across alternative definitions. These shares likely understate the extent of labelling because the pay movements used to construct them also contain changes in permanent earning power. Removing this component raises the shares to an upper bound of 0.176, 0.256, and 0.488.

These measured shares imply netting floors of 2.21, 3.63, and 3.74 $(\%)^2$ from (13). The difference between each floor and the estimated transitory prior, so the residual, therefore identifies the reading noise. It is essentially zero for the middle and top groups, but 7.43 $(\%)^2$ for the bottom group. The two sides of this calculation come from independent sources: the floors are measured from pay records, while the priors come from the SCE. The fact that the middle and top priors sit close to their netting floors is consistent with the mechanism's prediction that labelled pay lowers perceived transitory risk, while the excess at the bottom is consistent with the additional reading noise predicted for less legible pay.

5.5 Scoring the Estimates

I now ask whether the analytical results, evaluated at the measured labelled shares, can account for the estimated belief parameters. The filter predicts the updating gains, while Proposition 1 predicts the attribution shares.

Updating gains. The filter behind (8), evaluated at the measured shares, produces an implied updating gain for each income group. Greater legibility removes noise from the signal and should therefore allow the worker to update more quickly. The filter correctly predicts that the bottom updates most slowly, but understates the differences across groups: the implied gains are 0.76, 0.78, and 0.82 for the bottom, middle, and top, respectively, compared with estimated gains of 0.63, 0.91, and 0.95 (Appendix D.4). I therefore use the estimated gains in the quantitative model, while the filter provides a useful check on their ordering.

Attribution ratios. Proposition 1, evaluated at the measured shares, gives attribution shares of 0.677, 0.524, and 0.401, which decline with income as in Figure 4. I compare the mechanism to the estimated cross-group ratios rather than to the levels, since the cell-mean shock regressor attenuates the estimated responses.[5] The predicted bottom-to-top and middle-to-top ratios are 1.69 and 1.31, compared with estimated ratios of 3.04 and 1.76, with log standard errors of 0.63 and 0.59. Both predictions therefore lie within one standard error of the estimates. The measured delivery of pay explains roughly half of the decline in attribution with income, leaving the remainder to either additional under-reaction not captured by labelling or sampling uncertainty.

5.6 The Belief Profile

Since I let income drift continuously in the model, I interpolate the belief objects between the three estimated income groups and use the resulting profiles in the quantitative model. For the quantitative profile, I take the estimated updating gain and transitory prior as given. The attribution share and resting permanent-risk belief are instead generated from the measured labelled share using the closed-form expressions in Section 4. I interpolate the labelled share and updating gain in levels, and the prior in logs, using monotone interpolants through the three group estimates. Each object is held fixed beyond the observed income range. I also impose the netting floor from (13), censoring the prior below $(1 - \rho(y))\sigma^2_{\theta,g}(y)$. Appendix E.1 describes the interpolation and evaluates the resulting mapping at representative income levels. The resulting prior retains the dip at the middle because the netting floors are themselves non-monotone.

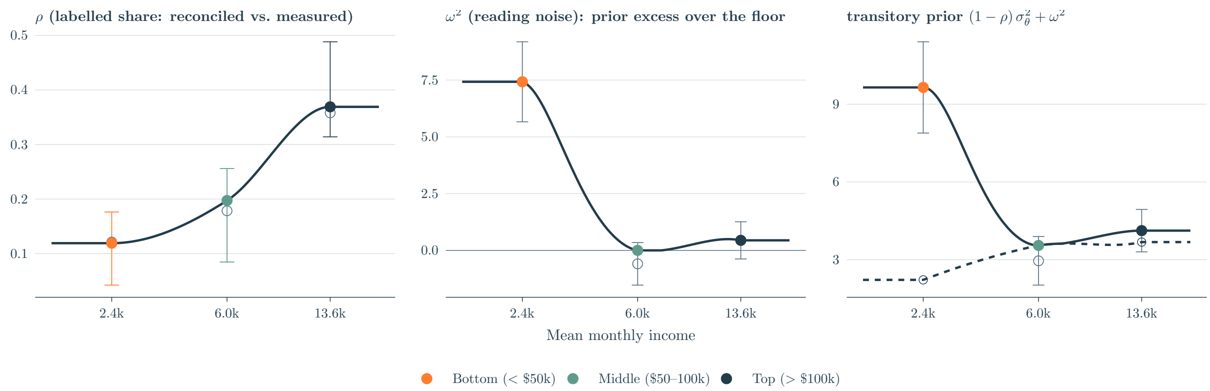

Figure 5: The labelled share ρ, the reading noise ω^2, and the transitory prior they compose (13), by income. Left: the share used in the model (filled points and income interpolant) against the measured share (open points), with whiskers spanning the specification range in Appendix D.2. Middle: reading noise, measured as the prior's excess over its netting floor, with ±1 standard-error bands in grey. Right: the carried prior, given by the netting floor plus reading noise, against the raw estimates and their ±1 standard-error bands in grey. In $(\%)^2$ except for the dimensionless share.

[5] Appendix D.3 estimates the corresponding group-specific loadings as 0.43, 0.56, and 0.58 for the bottom, middle, and top. Correcting for these loadings increases the estimated bottom-to-top ratio from 3.04 to about 4.1 and the middle-to-top ratio from 1.76 to 1.83, so the correction widens rather than eliminates the gradient. I retain the uncorrected ratios as the baseline because the loading correction is itself an approximation.

The resulting profile determines the permanent and transitory components of perceived risk at each income level. Figure 6 uses these profiles to construct the perceived-risk wedge and compare it with the full-information benchmark. We see the crossing conjectured by the mechanism: the permanent component remains below its realised counterpart, while the transitory component crosses the benchmark as income rises. The total perceived-risk wedge therefore also changes sign, with the crossing driven by the transitory component.

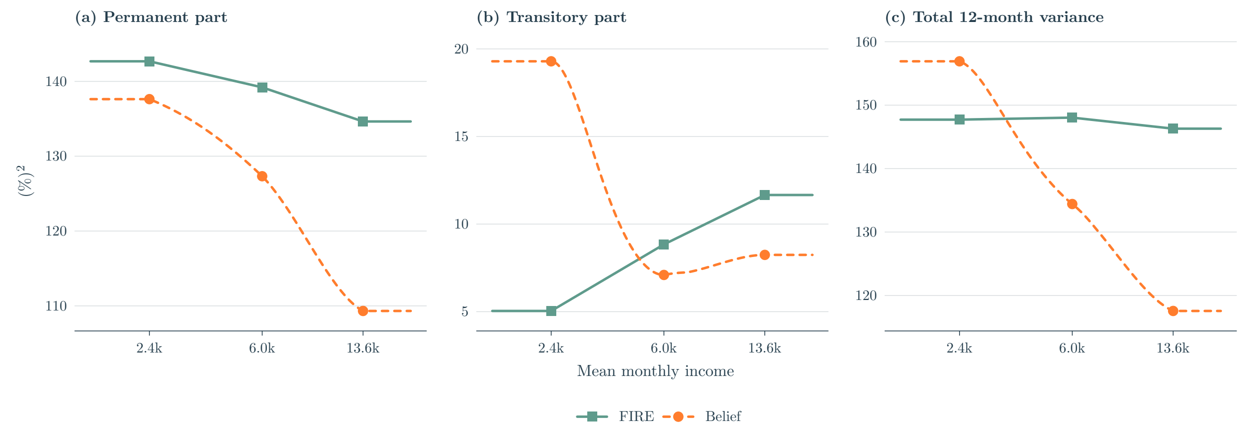

Figure 6: Perceived twelve-month risk against the FIRE benchmark, in closed form along the measured belief profile of Figure 5. Each panel draws the benchmark (solid) against its perceived counterpart (dashed), and the distance between the lines is the matching piece of the wedge (12): the permanent part $12\sigma^2_{\psi,g}$ against $12\hat{\sigma}^2_{\psi}$ in (a), the transitory part $2\sigma^2_{\theta,g}$ against $2\sigma^2_{\theta,\text{perc}}$ in (b), and the total in (c). Points mark the three income groups. In $(\%)^2$.

6 The Belief Wedge in a Heterogeneous-Agent Economy

I now embed the estimated belief structure in a general-equilibrium incomplete-market economy and ask how misperceived wage risk affects saving and consumption. Each income group is compared with an otherwise identical FIRE economy, such that the difference between the two regimes isolates the belief wedge. I first describe the environment and calibration, then quantify the wedge in the stationary cross-section and in general equilibrium, and finally study its dynamics when either the legibility of pay or the underlying wage risk changes.

6.1 Environment

Demographics and preferences. The economy is a monthly overlapping-generations life-cycle buffer-stock model following Wang (2023), closed with a single asset. Workers enter at age 25, work until 65, retire onto a pension replacing sixty percent of permanent income, and die with certainty at 85, with a small constant mortality risk before then. Each death is replaced by a new entrant drawn from the income distribution, so the stationary cross-section is maintained through finite lives and the permanent income component doesn't explode. Preferences are CRRA with $\gamma = 2$. At the terminal age, the household receives utility from its remaining assets through a bequest motive rather than being required to consume them fully.

Income. Labour income follows the permanent-transitory process estimated in Section 3, combined with a deterministic age profile: $y_{i,t} = \exp(p_{i,t})G_\tau \xi_{i,t}$, where $p_{i,t} = p_{i,t-1} + \psi_{i,t}$ is the permanent component and $\xi_{i,t} = \exp(\theta_{i,t})$ is the transitory component when employed. Their variances are taken from Table 2. The age profile G_τ is the hump-shaped deterministic component used by Wang (2023). When unemployed, the worker receives a benefit replacing one-half of permanent income; in retirement she receives the pension.

Mobility. Permanent innovations accumulate over the life cycle, so workers move through the income distribution even without changing jobs. A sequence of positive permanent innovations, for example, can move a worker from the bottom to the top of the distribution. The deterministic age profile adds predictable earnings growth but does not itself constitute wage risk. The income groups used to estimate beliefs are therefore interpreted in the model as positions in earning power rather than fixed worker types. Appendix C.6 checks the realised side of this mapping by re-sorting workers on their own earnings.

Unemployment. The wage-risk estimates condition on remaining employed, but unemployment is an important component of household income risk (Caplin et al., 2026). Following Wang (2023) and using his conversion, I therefore include a two-state employment process. Separation and job-finding rates are obtained from the SCE's perceived job-loss and job-finding probabilities.[6]

The resulting unemployment rates are 6.2/4.0/3.7 percent. These employment transitions are identical across the belief and FIRE regimes. They therefore affect the level and distribution of wealth, but are not themselves a source of the belief wedge. Beliefs update only while employed and remain fixed during unemployment.

Beliefs. The belief structure estimated in Section 5 enters through three income-dependent objects: the labelled share $\rho(y)$, the updating gain $\mu(y)$, and the perceived transitory prior $\sigma^2_{\theta,\text{perc}}(y)$. The labelled share is pinned within its measured range, while the updating gain and transitory prior are taken from the empirical estimates. The attribution share and resting permanent-risk belief are then obtained from the structural relationships in Section 4. As income changes, workers move continuously along these profiles.

Normalisation. Following Carroll (2006), I normalise all nominal quantities by permanent income. The state variable is therefore cash-on-hand per unit of permanent income, transitory income is $\exp(\theta)$, and permanent innovations enter through the growth factor $G_\tau \exp(\psi)$. The problem is scale invariant, so the group-specific mean monthly SCE incomes of \$2,415, \$6,014, and \$13,604 are used only to anchor magnitudes.

Technology and market clearing. The general equilibrium economy has a constant-returns production function $Y = ZK^\alpha N^{1-\alpha}$, with capital share $\alpha = 0.36$ and monthly depreciation implied by an annual depreciation rate of 2.5 percent. The single asset held by households is also the productive capital stock. Unemployment benefits and pensions are financed through balanced-budget labour-income taxes. The capital market clears at the equilibrium interest rate, while the productivity normalisation Z is fixed using the FIRE economy at a target annual interest rate of three percent and then held fixed across regimes.

[6] For each employed respondent, I take the reported probability of losing the main job over the next twelve months and, conditional on losing it, the probability of finding a new job within three months. I average these probabilities within income group using survey weights and convert a horizon-H probability P into a monthly hazard $-\log(1-P)/H$. The resulting monthly separation rates are 0.015/0.011/0.011 and job-finding rates are 0.23/0.27/0.29 from bottom to top.

Two regimes. I solve the economy under two regimes. In the *belief* regime, households form expectations using the estimated belief process. In the *FIRE* regime, they know the true permanent and transitory wage variances. The income process, employment transitions, preferences, technology, and fiscal system are otherwise identical. The difference between the two regimes therefore isolates the consequences of wage-risk perceptions.

6.2 The Household Problem

Let m denote cash-on-hand per unit of permanent income. The household chooses consumption subject to the borrowing constraint $c \leq m$:

$$\begin{aligned} V_\tau(m, p, s, \hat{\sigma}^2_\psi) &= \max_{0 \leq c \leq m} \left\{ u(c) + \beta\,(1 - D)\,\widetilde{\mathbb{E}}\left[\Gamma'^{1-\gamma}\, V_{\tau+1}(m', p', s', \hat{\sigma}^2_{\psi'})\right] \right\}, \\ m' &= \frac{R}{\Gamma'}\,(m - c) + \xi', \qquad \Gamma' = G_\tau \exp(\psi'). \end{aligned} \tag{18}$$

Here τ is age, p is the permanent-income state, s is employment status, D is the monthly mortality probability, and $R = 1 + r$ is the gross return on the asset. The expectation $\widetilde{\mathbb{E}}$ is taken under the household's perceived income process instead of the realised one. The terminal age is handled through the bequest motive described above.

The distinction between the budget constraint and the expectation is important. Households receive the realised income draw ξ' and permanent shock ψ', but choose saving using their perceived distribution of those shocks. Beliefs therefore affect behaviour through the precautionary-saving motive rather than through the realised budget constraint.

The permanent-risk belief evolves according to the constant-gain rule from Section 4,

$$\hat{\sigma}^2_{\psi,t} = \big(1 - \mu(y_t)\big)\hat{\sigma}^2_{\psi,t-1} + \mu(y_t)\lambda(y_t)(\Delta\tilde{y}_t)^2, \tag{19}$$

where $\Delta\tilde{y}_t$ is the worker's perceived, netted wage change. The attribution share $\lambda(y)$ is the structural value implied by the labelled share $\rho(y)$, while the updating gain $\mu(y)$ is taken from the empirical estimates. The transitory variance entering the perceived income distribution is the estimated, floor-censored $\sigma^2_{\theta,\text{perc}}(y)$.

For the stationary comparison, the belief state is fixed at its resting value from (11). The dynamic exercises instead carry $\hat{\sigma}^2_\psi$ as a state, allowing the model to trace how a sequence of wage shocks changes perceived risk and, through saving, the household's asset position.

Stationary equilibrium. For each regime, a stationary equilibrium consists of age-dependent value and policy functions and an invariant distribution Φ over household states such that households optimise given prices, Φ is invariant under the income, employment, demographic, and belief transitions, the capital market clears, and the government budget balances. Under FIRE, the belief state is degenerate at the true wage variances.

I compute stationary moments from a simulated panel of 30,000 households over 1,200 months, discarding the first 720 months as burn-in. The belief and FIRE economies are simulated using the same realised income and employment shocks, so differences between the two regimes are generated by beliefs rather than by different shock histories.

6.3 Calibration

Table 3 collects the parameters used in the quantitative exercise. The realised income and employment processes are taken from the estimates and measurements described above. For the belief block, the labelled share ρ_g is pinned within its measured range, while the updating gain μ_g and perceived transitory prior $\sigma^2_{\theta,\text{perc},g}$ are taken from the estimates of Section 5. The attribution share λ_g and resting permanent-risk belief $\hat{\sigma}^2_{\psi,g}$ are then implied by the structural relationships in Section 4. Thus the permanent belief is not separately calibrated.

The only remaining free preference parameter is the discount factor β, which I choose to match average liquid wealth in the FIRE economy. I use a common β across income groups so that the cross-sectional differences in saving behaviour arise from the income process, employment risk, and beliefs rather than from heterogeneous patience.

Parameter	Symbol	Value (Bot / Mid / Top)	Source or target
A. Preferences and assets (common)			
Risk aversion	γ	2	Assigned
Annual interest rate	r	2%	Fixed for PE; clears near 3% in GE
Borrowing limit	$\underline{a}$	0	Assigned
Discount factor	β	0.917	Calibrated to mean liquid wealth ($\approx$ 1.8 months)
B. Employment block (heterogeneous by income; common across regimes)			
Separation rate	s_g	0.015/0.011/0.011/mo	SCE perceived, monthly hazard
Job-finding rate	f_g	0.23/0.27/0.29/mo	SCE perceived, monthly hazard
Replacement rate	b	0.5	Assigned
Pension replacement	ς	0.6	Wang (2023)
C. Realised income process by group, $(\%)^2$			
Permanent variance	$\sigma^2_{\psi,g}$	11.9/11.6/11.2	SIPP (Table 2)
Transitory variance	$\sigma^2_{\theta,g}$	2.5/4.4/5.8	SIPP (Table 2)
Mean monthly income	–	\$2,415/6,014/13,604	SCE (scale anchor)
D. Beliefs by group			
Labelled share	ρ_g	0.119/0.197/0.369	Pinned within measured range
Update gain	μ_g	0.629/0.910/0.950	GMM estimate
Attribution	λ_g	0.677/0.519/0.398	Implied from ρ_g
Transitory prior	$\sigma^2_{\theta,\text{perc},g}$	9.65/3.55/4.12	GMM estimate, floor-censored
Resting permanent belief	$\hat{\sigma}^2_{\psi}$	11.47/10.61/9.11	Implied from belief rule

Table 3: Model calibration. The employment block is heterogeneous by income but common to both regimes. Under FIRE, households instead use the realised wage variances.

Discount factor. I calibrate the monthly discount factor β so that the FIRE economy, at an annual interest rate of 2 percent, matches average liquid wealth in the data, approximately 1.8 months of permanent income. This gives $\beta = 0.993$, equivalent to an annual discount factor of $\beta^{12} = 0.917$.[7] I hold this calibration fixed when comparing the belief and FIRE economies. Allowing patience to vary with income would introduce a second source of cross-sectional heterogeneity that is unrelated to the mechanism studied here.

[7]The target is constructed from the SCE Household Finance module using the liquid-wealth concept of Kaplan et al. (2014). Median liquid wealth is approximately 0.5/2.0/3.0 months of permanent income across the three income groups, giving an average close to 1.8 months.

6.4 The Cross-Sectional Wedge

I compare the stationary equilibrium under the estimated beliefs with an otherwise identical FIRE benchmark. The resulting differences in wealth and consumption responses isolate the behavioural consequences of the belief wedge.

Perceived risk. Table 12 reports stationary perceived twelve-month risk alongside FIRE. Perceived risk falls with income, from 153 to 139 to 124 $(\%)^2$ for the bottom, middle, and top, while FIRE is approximately 148 $(\%)^2$ across groups. The bottom therefore slightly over-perceives total risk, whereas the middle and top under-perceive it. This reflects the decomposition in Figure 6: the permanent component is increasingly discounted with income, while the transitory component crosses the FIRE benchmark.

Figure 7 compares the belief economy with FIRE. In the FIRE economy, mean wealth is 2.28, 1.71, and 1.50 months of permanent income across the bottom, middle, and top, with twelve-month MPCs of 0.470, 0.411, and 0.377; the full set of model moments is reported in Appendix F.

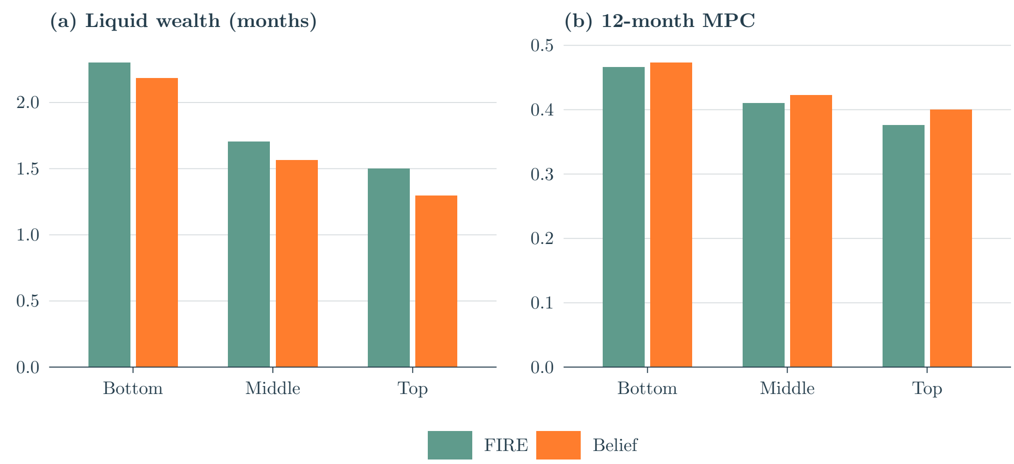

Figure 7: Mean liquid wealth, measured in months of permanent income, and the 12-month MPC by income group under FIRE and estimated beliefs.

The belief wedge in wealth increases monotonically with income. Relative to FIRE, the belief economy holds 5.1, 8.3, and 13.8 percent less wealth for the bottom, middle, and top, while the twelve-month MPC is higher by 0.7, 1.3, and 2.4 percentage points. The distortion therefore becomes substantially larger at higher incomes, consistent with the larger permanent-risk discount implied by the belief profile.

The bottom is affected least: its permanent belief is only modestly below the truth, while its higher unemployment risk also generates precautionary saving independently of wage-risk beliefs. The middle experiences a larger distortion, while the top responds most strongly because it combines the largest labelling share with the largest gap between perceived and realised permanent risk. The wedge operates only for households with slack above the borrowing constraint, and since hand-to-mouth households are already at the saving floor so therefore do not adjust saving in response to their perceived wage risk. The borrowing constraint thus does not drive the ordering. Restricting the analysis to workers and re-tercilling within workers gives wealth wedges of $-5.8/-9.7/-15.1$

percent and MPC differences of +0.9/+1.5/+2.7 percentage points.

To identify the source of the wedge, I solve two counterfactual economies. Holding the permanent belief at its FIRE value eliminates the wedge, whereas retaining the permanent belief and restoring the transitory prior reproduces it. The stationary behavioural wedge therefore operates through the perceived permanent component of wage risk; the transitory-risk misperception has little additional effect on saving or the MPC.

The result is robust to the interest rate. Recalibrating the discount factor to the same liquid-wealth target at the general-equilibrium rate of three percent leaves the wealth wedge at approximately 5.1, 8.3, and 13.6 percent. Holding the discount factor fixed instead raises it to 6.8, 10.5, and 16.4 percent. The corresponding general-equilibrium outcomes are reported in Appendix F.

6.5 The Closed Economy

Closing the economy allows the saving response to the belief wedge to feed back into the return to capital. I keep the productivity normaliser fixed across regimes and let the interest rate clear the capital market. The full results are reported in Appendix F.

The FIRE economy clears at the three-percent target by construction. Under the belief regime, aggregate capital is about five percent lower and the equilibrium return rises by 0.21 percentage points, from 3.00 to 3.21 percent. Lower perceived permanent risk therefore reduces precautionary saving and raises the return, consistent with Wang (2023). The aggregate response is modest relative to the cross-sectional wedge: the main effect remains the larger wealth and MPC differences across income groups.

This result reflects the one-asset closure, in which the same asset serves as both the household's precautionary buffer and the economy's capital stock. A richer treatment with separate liquid and illiquid assets could allow the belief wedge to affect portfolio composition as well as aggregate saving.

6.6 Dynamics

I now trace the belief wedge through three experiments motivated by the comparative statics in Section 4. The first changes how legibly pay is delivered by raising the labelled share of low-paid work. The second raises the underlying variance of permanent income risk. The final experiment is a recession, which changes neither primitive. In each case I hold prices fixed and compare the belief economy with an otherwise identical FIRE economy.

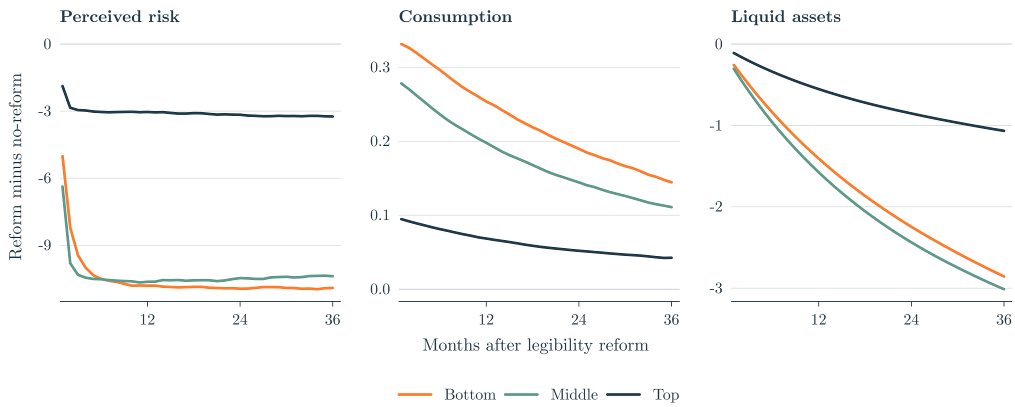

Figure 8: Response to a legibility reform that raises the labelled share of the bottom and middle to the top's level, $\rho = 0.37$. The belief regime is shown relative to the no-reform path. Perceived risk is the twelve-month perceived variance in $(\%)^2$, while consumption and liquid assets are in percent of the no-reform path. Paths average twelve simulated cohorts.

A legibility shock. The first experiment raises the labelled share of low-paid workers to the top group's level, leaving the realised income process unchanged. Under the estimated belief rule, greater labelling reduces the share of a wage surprise attributed to permanent risk. Perceived permanent risk therefore falls, precautionary saving declines, and consumption rises. The effect is largest for the bottom and middle groups, whose labelled shares increase most: liquid wealth falls by about 2.9 and 3.0 percent, respectively, compared with 1.1 percent for the top. The belief adjusts within months, while the wealth response unfolds over one to two years. Since the FIRE economy is unaffected by the reform, these responses isolate the belief channel.

Figure 9: Response to a persistent 40% increase in true permanent wage variance, belief regime relative to FIRE. Perceived risk is the change in twelve-month variance in $(\%)^2$, while consumption and liquid assets are in percent of the no-shock baseline. Paths average twenty-four simulated cohorts.

A risk shock. The second experiment raises the true variance of permanent income risk by 40% and holds the increase persistent. Under FIRE, households immediately recognise the higher risk and rebuild their precautionary buffers. Under the belief rule, households attribute only part of the additional squared wage surprises to permanent risk. Consequently, the increase in perceived risk is smaller and the accumulation of liquid wealth is weaker. After two years, the belief economy has accumulated 14.7/18.1/18.7 percent less additional liquid wealth than FIRE for the bottom, middle, and top groups, respectively. The experiment therefore shows that the same under-reaction that generates the steady-state wedge also attenuates the precautionary response to an increase in underlying risk.

Finally, consider a recession represented by a one-off three-percent fall in income. This changes the level of earnings but neither the labelled share nor the underlying variance of permanent risk. The shock therefore produces only a temporary movement in the belief, with no persistent difference in precautionary wealth between the belief and FIRE economies. The wedge is thus tied to the primitives governing how risk is perceived, rather than to temporary movements in income alone.

7 Conclusion

This paper asks how much risk households believe they face in their own earnings, and what follows when those beliefs differ from the risk realised in the data. Using the SCE density forecasts and a matched same-job panel from the SIPP, I find that perceived wage risk varies substantially with income even though realised wage risk does not. The difference is not confined to the level of perceived risk: it appears in the persistence of beliefs and in their response to realised wage shocks.

I show that the income gradient has two sides. Low-income workers over-perceive transitory risk, while middle- and high-income workers under-perceive permanent risk. I interpret this pattern through how pay is delivered. A larger labelled share of transitory pay makes a wage change easier to net out and reduces the weight placed on it when learning about permanent earning power. At lower incomes, where pay is less clearly separated into permanent and transitory components, additional reading noise can instead raise perceived transitory risk. The estimated updating rule and the labelled share measured directly from pay records account for a substantial part of the observed gradients.

Embedding these beliefs in a heterogeneous-agent life-cycle economy gives the wedge economic significance. Under the measured belief profile, households hold less liquid wealth than under full information, with the reduction increasing with income, and the MPC wedge is correspondingly concentrated at the top. In general equilibrium, the aggregate response is modest: lower precautionary saving reduces the capital stock and raises the equilibrium return, consistent with the mechanism identified by Wang (2023). The dynamic experiments further show that the same belief mechanism affects how households respond to changes in the delivery and riskiness of earnings.

The broader implication, and what the practitioner should take away, is that income risk in incomplete-market models is not only a property of the earnings process but also a property of how that process is perceived. Measuring those perceptions, and the way households form them, can therefore change both the cross-section of precautionary saving and the response of households to changes in income risk.

References

Augenblick, N., Lazarus, E., and Thaler, M. (2025). Overinference from weak signals and underinference from strong signals. *Quarterly Journal of Economics*, 140(1):335–401.

Ben-David, I., Fermand, E., Kuhnen, C. M., and Li, G. (2018). Expectations uncertainty and household economic behavior. Working Paper 25336, National Bureau of Economic Research.

Blundell, R. and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1):115–143.

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.

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

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.

Caplin, A., Gregory, V., Lee, E., Leth-Petersen, S., and Sæverud, J. (2026). Subjective earnings risk. *Review of Economic Studies*. Advance access, doi:10.1093/restud/rdag070.

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

Carroll, C. D., Crawley, E., Slacalek, J., Tokuoka, K., and White, M. N. (2020). Sticky expectations and consumption dynamics. *American Economic Journal: Macroeconomics*, 12(3):40–76.

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

Druedahl, J. and Jørgensen, T. H. (2020). Can consumers distinguish persistent from transitory income shocks? *Economic Journal*, 130(632):2410–2437.

Edwards, W. (1968). Conservatism in human information processing. In Kleinmuntz, B., editor, *Formal Representation of Human Judgment*, pages 17–52. Wiley, New York.

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

Enke, B. and Graeber, T. (2023). Cognitive uncertainty. *Quarterly Journal of Economics*, 138(4):2021–2067.

Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.

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.

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

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.

Koşar, G. and Melcangi, D. (2025). Subjective uncertainty and the marginal propensity to consume. Staff Report 1148, Federal Reserve Bank of New York.

Koşar, G. and van der Klaauw, W. (2025). Workers' perceptions of earnings growth and employment risk. *Journal of Labor Economics*, 43(S1):83–121.

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.

Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, 126(1):373–416.

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

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.

Muth, J. F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association*, 55(290):299–306.

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

Pistaferri, L. (2001). Superior information, income shocks, and the permanent income hypothesis. *Review of Economics and Statistics*, 83(3):465–476.

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

Sanchez, J. M. and Song, S. (2025). Earnings misperceptions and household distress. Working Paper 2025-030, Federal Reserve Bank of St. Louis.

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

A 12-Month FIRE Variance

Rolling $e_t = p_t + \theta_t$ forward H months, the permanent piece accumulates one fresh innovation per month, $\text{Var}(p_{t+H} - p_t) = H\sigma^2_{\psi,g}$, while the iid transitory piece contributes only its endpoints, $\text{Var}(\theta_{t+H} - \theta_t) = 2\sigma^2_{\theta,g}$. Since the shocks are independent, $\text{Var}(e_{t+H} - e_t) = H\sigma^2_{\psi,g} + 2\sigma^2_{\theta,g}$, which at $H = 12$ gives (5).

Group	FIRE benchmark ($12\sigma^2_{\psi,g} + 2\sigma^2_{\theta,g}$)	SCE observed $\bar{\sigma}^2_{\text{perc}}$	ratio
Bot	147.7	20.6	7.2×
Mid	148.0	11.3	13.1×
Top	146.3	11.4	12.8×

Table 4: FIRE-implied 12-month variance versus SCE-reported perceived variance, in $(\%)^2$.

Appendix C.7 recomputes the benchmark under the same-hours restriction. It falls to 79, 85, and 98 $(\%)^2$ for the bottom, middle, and top groups, respectively, while the corresponding ratios fall to 3.8×, 7.5×, and 8.6×.

B Closed Form for λ_g under Partial Labelling

This appendix derives the closed form (10). The attribution on the full pay change is the unbiased attribution on the unlabelled residual, attenuated by the labelled pay that enters the change. Split the transitory shock $\theta_t = \theta^l_t + \theta^u_t$ into independent pieces with variances $\rho_g\sigma^2_{\theta,g}$ and $(1-\rho_g)\sigma^2_{\theta,g}$. Then $\Delta e_t = u_t + l_t$, where the unlabelled residual is $u_t = \psi_t + \theta^u_t - \theta^u_{t-1}$ and the labelled component is $l_t = \theta^l_t - \theta^l_{t-1}$. Hence, $V_{u,g} = \sigma^2_{\psi,g} + 2(1-\rho_g)\sigma^2_{\theta,g}$, $V_{l,g} = 2\rho_g\sigma^2_{\theta,g}$, and $V_{u,g} + V_{l,g} = \sigma^2_{\psi,g} + 2\sigma^2_{\theta,g} = \bar{S}_{m,g}$, which is the variance of the full monthly change.

The rule of Section 4 updates on the full observed change $(u+l)^2$. The attribution recovered by (16) is therefore the unbiased residual attribution $\sigma^2_{\psi,g}/V_{u,g}$ attenuated by the projection of u^2 on $(u+l)^2$, so $\lambda_g = (\sigma^2_{\psi,g}/V_{u,g})\,\text{Cov}(u^2,(u+l)^2)/\text{Var}((u+l)^2)$. Under joint Gaussianity with $u \perp l$, $\text{Cov}(u^2,(u+l)^2) = 2V^2_{u,g}$ and $\text{Var}((u+l)^2) = 2(V_{u,g} + V_{l,g})^2$. It follows that

$$\lambda_g = \frac{\sigma^2_{\psi,g} V_{u,g}}{(V_{u,g} + V_{l,g})^2} = \frac{\sigma^2_{\psi,g}\left(\sigma^2_{\psi,g} + 2(1-\rho_g)\sigma^2_{\theta,g}\right)}{(\sigma^2_{\psi,g} + 2\sigma^2_{\theta,g})^2},$$

which is (10).

This closed form is one of three readings of the estimating equation, and the scorecard's ratio comparison does not depend on which reading is used. Under the netting reading, in which the agent nets first and updates on u^2_t at the fixed attribution $\lambda^*_g = \sigma^2_{\psi,g}/\bar{S}_{m,g}$, the coefficient in (16) is the projection $\lambda^*_g\,(V_{u,g}/\bar{S}_{m,g})^2 = \sigma^2_{\psi,g}V^2_{u,g}/\bar{S}^3_{m,g}$. The predicted bottom-to-top and middle-to-top ratios become 1.99 and 1.48, respectively, instead of 1.69 and 1.31. Both pairs lie within one standard error of the estimated ratios (Appendix D.3). The implemented form, which updates on the agent's own noisy netted reading with the rescaled attribution $\tilde{\lambda}$ (Appendix D.1), shares its resting mean with both alternatives. Thus, the readings differ in their coefficient rather than in their resting belief.

The Gaussian fourth-moment identities can also be checked directly on the component panel of Appendix D.2. Replacing them with the sample projection $\mathrm{Cov}(u^2, (u+l)^2)/\mathrm{Var}((u+l)^2)$ moves the predicted attribution ratios from 1.69 and 1.31 to 1.54 and 1.13 under the baseline trim, still about one standard error from the estimated ratios. The untrimmed projection, by contrast, is dominated by the labelled bonus tail documented there.

C Belief Estimates and Robustness

Table 5 reports the headline joint-GMM estimates of (16) plotted in Figure 4. Sections C.1–C.6 assess their robustness, C.7 constructs the same-hours counterpart of the realised process, and C.8 measures the twelve-month benchmark directly.

Group	$\sigma^2_{\theta,g}$ (SIPP truth)	$\sigma^2_{\theta,\mathrm{perc},g}$ (GMM)	λ_g (GMM)	μ_g (GMM)	Hansen J p	N
Bot	2.52	**9.65** (1.76)	+**0.079** (0.031)	0.629 (0.055)	0.23	1,531
Mid	4.42	2.96 (0.94)	+0.046 (0.015)	0.910 (0.089)	0.23	7,501
Top	5.83	4.12 (0.82)	+0.026 (0.013)	0.950 (0.060)	0.23	11,331
Wald ($B = T$)	—	$p =$ **0.004**	$p = 0.12$	$p =$ **0.0001**	—	—

Table 5: Joint Blundell-Bond system GMM on (16), by income group (the estimates of Figure 4). Two-way clustered SEs (respondent, quarter) in parentheses. Instruments: $\bar{\sigma}^2_{\mathrm{perc}}$ at lags 3–4, the shock at lags 1–3, plus BB level moments, collapsed. N counts observations, and respondent counts are 269/1,050/1,452. The Hansen J is a single joint-system statistic repeated across rows. Bold p-values mark Wald rejections of cross-group equality at 5%, and bold estimates mark the bottom's prior and attribution (Section 5).

C.1 Occupation × Time Fixed Effect

Table 6 adds an occupation × year-month FE to the shock construction. The persistence and attribution gradients survive, attribution becomes more precise, and the prior gradient retains its sign but loses significance.

	Canonical			Occupation × time FE		
Group	$\hat{\mu}_g$	$\hat{\lambda}_g$	$\hat{\sigma}^2_{\theta,\mathrm{perc},g}$	$\hat{\mu}_g$	$\hat{\lambda}_g$	$\hat{\sigma}^2_{\theta,\mathrm{perc},g}$
Bot	0.629 (0.055)	+**0.079** (0.031)	**9.65** (1.76)	0.649 (0.061)	+**0.088** (0.039)	**7.30** (2.66)
Mid	0.910 (0.089)	+0.046 (0.015)	2.96 (0.94)	0.923 (0.090)	+0.037 (0.010)	2.62 (0.98)
Top	0.950 (0.060)	+0.026 (0.013)	4.12 (0.82)	0.959 (0.060)	+0.017 (0.008)	4.19 (0.73)
Wald μ (T−B)		$p =$ **0.0001**			$p =$ **0.0003**	
Wald λ (B−T)		$p = 0.116$			$p = 0.068$	
Wald σ^2 (B−T)		$p =$ **0.004**			$p = 0.260$	

Table 6: Canonical versus occupation × year-month FE in the wage-shock construction. Standard errors in parentheses.

C.2 Lag Depth, Trim, COVID, and the Bottom's Trend

Specification variants. Table 7 varies the lags, trim, and COVID handling one at a time. The Hansen J never rejects, attribution remains monotone in every specification, the persistence Wald survives all but the tighter trim, and the prior Wald survives all but the COVID exclusion.

	$\hat{\mu}_g$			$\hat{\lambda}_g$			$\hat{\sigma}^2_{\theta,\text{perc},g}$			Wald p (B vs T)		
Variant	B	M	T	B	M	T	B	M	T	μ	λ	σ^2
Canonical	0.629	0.910	0.950	+0.079	+0.046	+0.026	9.65	2.96	4.12	**0.0001**	0.116	**0.004**
Trim 2.5/97.5	0.824	0.846	0.891	+0.096	+0.055	+0.014	11.16	4.05	4.82	0.067	0.188	**0.000**
$\bar{\sigma}^2$ lags 2–3	0.649	0.951	0.876	+0.077	+0.050	+0.020	9.85	2.87	4.44	**0.000**	0.080	**0.004**
$\bar{\sigma}^2$ lags 4–5	0.590	0.899	0.948	+0.104	+0.046	+0.021	8.94	3.01	4.38	**0.000**	0.063	**0.040**
Shock lags 0–1	0.681	0.936	0.940	+0.084	+0.016	+0.012	9.67	4.42	4.89	**0.013**	0.061	**0.012**
Shock lags 2–4	0.672	0.908	0.948	+0.098	+0.040	+0.027	8.67	3.21	4.10	**0.002**	**0.010**	**0.016**
Exclude 2020–21	0.642	0.992	0.886	+0.140	+0.043	+0.013	6.40	3.13	4.88	**0.021**	0.052	0.637

Table 7: One-at-a-time variations around the canonical specification. Bold p-values reject cross-group equality at the 5% level.

The bottom's trend. The bottom series drifts upwards (+0.83 per year, $t = 3.7$), which puts some pressure on the mean-stationarity assumption underlying the Blundell-Bond level moments, although the belief block does not rely on those moments. A linear detrend or calendar-year effects leave the bottom prior at 9.74 and 9.57, respectively, against the canonical 9.65, and the bottom gain at 0.62 and 0.64 against 0.63. The gain Wald remains below 0.0001, while the prior Wald is at most 0.008. The difference-in-Hansen statistic for the level moments does not reject ($p = 0.55$ and 0.44 at the bottom and top), and dropping them preserves the point estimates, with the bottom prior at 9.22 (1.81), at the cost of some precision in the gain.

C.3 The Sign of News

Positive wage news carries roughly twice the belief weight of negative news. I treat this asymmetry as suggestive rather than established. Pooling the groups for power and splitting the shock by sign, both components load positively, with $\lambda^+ = 0.055$ versus $\lambda^- = 0.027$, significant across cluster choices ($N = 20{,}363$, Hansen J $p = 0.11$). The direction is opposite to what loss aversion would predict (Kahneman and Tversky, 1979): upside news appears to be treated as informative about future earning power, while downside news is more often treated as transitory. Two caveats make the interpretation suggestive. First, the additional weight may reflect the size of the surprise rather than its sign, since expected wage growth was positive over the sample. Second, pooling leaves open whether the asymmetry itself differs by income.

C.4 Numeracy and Inflation Uncertainty

The level gradient of Section 3 survives conditioning on numeracy and inflation uncertainty. The SCE's numeracy module classifies 47.0 percent of bottom-group respondents as low-numeracy, compared with 28.4 percent in the middle and 11.1 percent at the top. On the full matched panel, which is broader than the estimation panel, the bottom-to-top difference in perceived variance is 10.6 (1.0). It falls to 2.4 (0.8, $p = 0.002$) among high-numeracy respondents and rises to 17.0 (2.1) among low-numeracy respondents.

Inflation uncertainty is itself strongly graded, at 27.9/15.2/8.5 across the groups, and controlling for it absorbs about seventy percent of the bottom-to-top wage-risk gap, leaving a wage-specific gradient of +3.09 (0.60). Under the over-control reading of Section 3, both controls remove part of the reading-noise channel itself, so the conditioned gaps are lower bounds. Under the alternative interpretation as mechanical reporting width, they are the cleaner estimates. The gradient survives under either interpretation, with its magnitude bracketed by the two.

Re-estimating the belief block (16) within high-numeracy respondents reproduces the middle and top, with priors of 2.35 (0.85) and 4.93 (1.02), respectively, at or below their realised values as in the baseline. The bottom retains only 75 estimation observations under the numeracy-module overlap, which is too few for system GMM, so its prior excess cannot be checked at the estimation level.

Three checks bound the extent to which reporting mechanics can account for the gradient. First, bottom respondents use 5.7 of the ten histogram bins, compared with 4.9 in the middle and 4.6 at the top. Restricting the sample to reports that use three or more bins leaves the gradient at 23.0/12.3/12.2. Second, low numeracy widens the reported inflation density by more than the wage density in every group, with wage-against-inflation log-variance gaps of 0.92 against 1.17 at the bottom, 0.47 against 0.76 in the middle, and 0.15 against 0.62 at the top. The widening is therefore a generic reporting trait rather than a feature specific to wage risk.

Third, under the survey's density-implied variance, which is a parametric fit in the tradition of Engelberg et al. (2009), the group means are 28.6/16.4/11.3, compared with the midpoint measure's 20.8/11.2/11.3 on the same reports from the broader panel. The gradient is therefore steeper under the density-implied measure. This qualifies Section 3: the middle-top equality there is a property of the midpoint measure, while both the ordering and the bottom's excess survive under either measure.

C.5 Persistence of the Realised Volatility Process

Realised volatility is more persistent at the top than at the bottom. The belief-persistence gradient in Table 1 therefore runs opposite to the realised process it would have to inherit. I estimate the realised counterpart using the same estimator as the belief AR(1): a Blundell-Bond system GMM ARCH(1) in the squared wage change, with lags 3–4 and two-way clustered standard errors. Thus, the two sides differ in what they measure, but not in how they are estimated.

	Monthly		Quarterly	
Group	Perceived $\hat{\phi}_g$	Realised $\hat{\alpha}_{1,g}$	Perceived $\hat{\phi}_g$	Realised $\hat{\alpha}_{1,g}$
Bot	+0.343 (0.046)	+0.181 (0.014)	+0.632 (0.062)	+0.010 (0.018)
Mid	+0.142 (0.026)	+0.196 (0.011)	+0.594 (0.088)	+0.028 (0.020)
Top	+0.144 (0.025)	+0.234 (0.013)	+0.377 (0.076)	+0.058 (0.016)
Wald ($B = T$)	$p = 0.0001$	$p = 0.006$	$p = 0.009$	$p = 0.050$

Table 8: Perceived-variance AR(1) persistence (SCE) against realised wage-volatility ARCH(1) persistence (SIPP), by income group. Both come from Blundell-Bond system GMM on the same instrument set, lags 3–4 in differences plus the BB level moment, collapsed, with two-way clustered standard errors in parentheses. The realised series is stripped of the negative first-order autocorrelation that differencing mechanically induces. Respondents number 732/1,384/1,676 on the SCE side and 28,262/43,743/55,162 on the SIPP side at monthly cadence.

The monthly realised gradient rises from 0.181 at the bottom to 0.234 at the top and rejects equality at $p = 0.006$. By contrast, the perceived gradient falls from 0.343 to 0.144 and rejects equality at $p = 0.0001$. The quarterly comparison points in the same direction, although the realised coefficients are small and individually insignificant below the top.

A cell-level check on the raw series underlying the shock regressor gives the same qualitative result. The within-cell monthly AR(1) is 0.29 at the bottom against 0.36 at the top, and 0.44 against 0.49 in the national series. Instrumenting the lag to remove finite-cell sampling noise drives the bottom to 0.06 against 0.26 at the top. The uninstrumented numbers therefore flatter the bottom, if anything, since adjacent squared changes share a mechanical correlation through the same person-month, and that correlation is larger where cells are smaller, as they are at the bottom.

C.6 The Realised Process: Higher-Order Moments and the Stayer Conditioning

The higher-order autocovariances of the wage change are economically second order, so the permanent-transitory model behind Table 2 provides an adequate description of the realised process. The model implies $\mathrm{Cov}(\Delta e_t, \Delta e_{t-k}) = 0$ for all $k \geq 2$. At these sample sizes, the estimated moments are statistically distinguishable from zero but remain small, between 2 and 23 percent of the lag-1 moment that identifies $\sigma^2_{\theta,g}$. The largest deviations occur at the quarterly lags for the top, +0.99 and +0.41 at lags three and six, consistent with the quarterly bonus-delivery mechanism in Section 5.

The stayer conditioning does not select the sample in a way that could understate the bottom's transitory risk. Months immediately preceding an employer change carry elevated volatility at the bottom, about 1.6 times that in other months, but they account for only one percent of the sample. Excluding them therefore moves $\sigma^2_{\theta,g}$ by less than 0.04 $(\%)^2$ in every group. Closing the bottom's prior gap through selective separation would instead require the conditioning to hide a fourfold understatement of its transitory risk, far beyond the observed elevation.

The sorting choice also does not drive the realised-process results. Sorting workers by their own annual earnings rather than household income reassigns about half of them, since spousal income moves many workers across the household cut-offs, but the pattern of Table 2 remains intact. The transitory variance rises more steeply with income, at 2.28, 3.49, and 8.11 $(\%)^2$ across own-earnings terciles, while the permanent variance remains roughly flat, at 12.44, 11.30, and 10.86. The realised gradient therefore runs opposite to the perceived gradient under either sort. The belief side cannot be re-sorted, because the SCE elicits household income alone, so the comparison across surveys retains the household sort.

C.7 The Same-Hours Counterpart of the Realised Process

The SCE question holds hours fixed, while the recorded SIPP wage carries them: earnings are divided by the average weekly hours reported for the month, premium pay and tips enter the numerator, and the reported schedule sets the divisor. The relevant counterpart therefore restricts wage changes to months in which reported hours do not change. I recompute the moments of Table 2 using only consecutive within-spell changes across which reported average weekly hours are unchanged, requiring both differences of each autocovariance pair to qualify and applying the full-sample trim bounds. Thus, the restriction rather than the trim drives the difference in the estimates. Within the restricted pairs, the wage change is the change in earnings at a fixed divisor, which also removes the division bias present in the baseline.

The restricted moments should be interpreted as an upper bound on the requested object rather than as a direct measure of it. Reported hours are heaped, so unchanged reports can conceal genuine changes, and hours worked may still vary around an unchanged reported value.

The restriction is mild, dropping 2.3, 1.8, and 1.4 percent of observations from the bottom to the top. Selection on the realised outcome therefore has little room to operate, although the dropped share itself falls with income, providing another margin of the exposure gradient. The dropped observations are nevertheless economically important: they account for thirty-one percent of total monthly variance at the bottom against sixteen percent at the top. A placebo that drops the same share of changes at random leaves the permanent variance within 0.13 $(\%)^2$ of the baseline in every group. It is therefore the months removed, rather than the number of observations lost, that drive the change.

	Baseline		Same reported hours		
Group	$\sigma^2_{\theta,g}$	$\sigma^2_{\psi,g}$	$\sigma^2_{\theta,g}$	$\sigma^2_{\psi,g}$	Share kept
Bot	2.52	11.92	2.83	6.12	0.977
Mid	4.42	11.61	4.73	6.31	0.982
Top	5.83	11.23	6.07	7.13	0.986

Table 9: SIPP permanent-transitory identification with and without the same-hours restriction, monthly cadence, in $(\%)^2$. The baseline columns re-estimate the construction of Table 2 on the same code path, matching its moments to within 0.05; the same-hours columns keep only within-spell changes across which the reported average weekly hours are unchanged, with both differences of each autocovariance pair required to qualify. Share kept is the fraction of baseline observations surviving the restriction.

Table 9 sets the restricted moments against the baseline. The transitory variance is essentially unchanged and still rises with income, so the realised gradient runs opposite to the perceived gradient under either conditioning. The netting floors of Appendix D.3 move to 2.49, 3.89, and 3.90, while the middle and top priors remain within one standard error of them. The bottom's excess is 7.16 ($t = 4.1$), compared with 7.43 at baseline. Rescaling the labelled share to the restricted movement, which removes hours-event variance from its denominator, raises the shares to about 0.17, 0.23, and 0.43 and changes none of the three comparisons. The full-information attribution ratio $\lambda^*_{\text{Bot}}/\lambda^*_{\text{Top}}$ moves from 1.43 to 1.40, while the closed form (10) at the restricted moments predicts attribution ratios of 1.71 and 1.25, instead of 1.69 and 1.31. These remain within one standard error of the estimates. The persistence comparison of C.5 is equally insensitive: it shows no bottom-top gap whether the wage is measured per reported hour or in raw terms.

The permanent variance falls by roughly half at the bottom and middle and by over a third at the top. A change in the reported schedule is rare but persistent. It is a move between full- and part-time work or a renegotiated contract rather than a noisy month, so the moments of Section 3 assign what it carries to the permanent component. The hours-linked permanent share, 5.8, 5.3, and 4.1 $(\%)^2$ from the bottom to the top, itself falls with income, showing the delivery gradient of

Section 5 on another margin. The baseline therefore retains this component because a persistent change in scheduled hours is precisely the kind of earning-power risk the household absorbs.

The twelve-month benchmark (5) falls to 79, 85, and 98 $(\%)^2$ and now rises with income where the reported variance falls. The correction therefore does not cancel from the cross-group comparison; it strengthens the contrast. Reported variance remains about four to nine times below the tighter benchmark in every group, so the shortfall on which the constant-attribution evidence of Section 4 rests is preserved.

The removed variance also gives the bottom's prior excess a second possible source. A bottom worker unable to comply fully with the instruction, who folded part of the hours-event risk into her answer, would be reporting truthfully and would still sit above her netting floor. Folding in about a fifth of the removed variance would account for the excess in full. The moments cannot distinguish this partial inclusion from unresolved reading of the pay she does condition on. At the same time, the inflation-density widening of Appendix C.4, which no hours channel can generate, indicates that part of the excess is generic reporting rather than either mechanism. Little turns on the allocation for behaviour, since each interpretation enters the transitory prior rather than the permanent belief. The excess is therefore best described as reported risk beyond the object the question asks about, with reading noise being the component the mechanism predicts rather than the entirety of what the data establish.

The baseline moments remain those of Table 2, because the budget of Section 6 must carry the earning-power risk the household actually absorbs, and a persistent cut in scheduled hours is precisely such a risk. The question conditions it out of the elicited object, not out of her income. The cost of this choice is bounded in closed form: at the restricted moments and measured shares, the resting discounts $V_u/\bar{S}_m$ deepen to 0.94, 0.89, and 0.77, while the twelve-month permanent gaps shrink from −5.0, −11.9, and −25.3 to −4.3, −8.1, and −19.3 $(\%)^2$. The wedge of Section 6 therefore retains its shape and its concentration at the top, at magnitudes smaller by between a seventh and a third.

C.8 The Benchmark at Long Horizons

The twelve-month benchmark of Section 3 extrapolates monthly moments under the random walk rather than measuring the twelve-month change directly. This subsection therefore measures the latter on the same panel. For each horizon h, I take within-spell pairs h months apart, drop pairs whose later endpoint is January as in the monthly construction, and trim the squared change at the same one and ninety-nine percent cutoffs. At $h = 1$, the exercise returns the moments of Table 2 by construction, 16.9, 20.4, and 22.9 $(\%)^2$. At $h = 12$, it returns 2724, 2610, and 2774, about eighteen times the extrapolated benchmark.

The gap is a statement about outliers rather than necessarily about the underlying wage process. A job stayer's monthly change is concentrated at zero with heavy tails: most months carry no change in the reported rate, so the mean-based moment is governed by the tail, and the common trim removes about eighty-five percent of it at one month against under half at twelve months. Measuring the same profile with a scale that ignores the tails, the interquartile range converted into Gaussian variance units, moves the comparison in the other direction: it puts the monthly variance near 0.6, far below the trimmed moment, and the twelve-month variance near 340, twice the benchmark rather than eighteen times it. One measure is governed by the tail and the other by the spike, and the data alone do not discipline the choice between them.

The level of the benchmark is therefore not identified without a model of the reporting outliers, which is why Section 3 reads it as an upper bound and bases the argument on the cross-group contrast rather than the level. What the profile does establish is the direction: the directly measured twelve-month variance lies above the extrapolation under either scale. The shortfall of reported risk relative to realised risk is therefore not an artefact of the random-walk extrapolation.

D Identification: the Filter, the Labelled Share, and the Scorecard

This appendix describes the filter behind the gain (D.1), the measurement and robustness of the labelled share (D.2), the scorecard that compares the measured mechanism with the estimated belief moments (D.3), the filter's diagnostic on the estimated gains (D.4), and the hours evidence underlying the reading-noise gradient (D.5).

D.1 The Filter

Once the attribution is rescaled to the agent's own signal, the perceived transitory prior drops out of the filter's signal-to-noise ratio. The gain therefore depends on the labelled share rather than on reading noise.

The agent's monthly signal is her own squared netted reading, $x_t = (\Delta\tilde{y}_t)^2$ with $\tilde{y}$ from (7). It has mean $\sigma^2_{\psi,g} + V_{x,g}$ and, under Gaussianity, variance $2(\sigma^2_{\psi,g} + V_{x,g})^2$, where $V_{x,g} = 2\sigma^2_{\theta,\mathrm{perc},g}$ is twice the perceived transitory prior. The perceived permanent variance follows a random walk with innovation variance q (the drift), which is the filter's one free parameter and is fitted to the estimated gains in D.4. The observation is the attributed signal $z_t = \tilde{\lambda}_g\, x_t = \hat{\sigma}^2_{\psi,t} + \varepsilon_t$, a unit-loading reading of the state, where $\tilde{\lambda}_g = \sigma^2_{\psi,g} V_{u,g} / [\bar{S}_{m,g}(\sigma^2_{\psi,g} + V_{x,g})]$ is the full-change attribution (10) rescaled to the signal's mean and $V_{u,g} = \sigma^2_{\psi,g} + 2(1-\rho_g)\sigma^2_{\theta,g}$ is the unlabelled residual variance of the main text. The rescaling makes the observation unbiased at rest, since $\tilde{\lambda}_g\, \mathbb{E}[x_t] = \sigma^2_{\psi,g}\, V_{u,g}/\bar{S}_{m,g}$ is the permanent belief (11). The variance of a squared signal moves with the state, so I evaluate it at the resting point and interpret the filter as moment-matched rather than exact. Its steady-state gain is $\mu_g = \bar{P}_g/(\bar{P}_g + R_g)$, with $\bar{P}_g = \frac{1}{2}\left(q + \sqrt{q^2 + 4qR_g}\right)$, where $\bar{P}_g$ is the positive root of the Riccati equation. The observation noise is $R_g = \mathrm{Var}(z_t) = \left(\tilde{\lambda}_g\right)^2 2(\sigma^2_{\psi,g} + V_{x,g})^2 = 2(\sigma^2_{\psi,g})^2\,(V_{u,g}/\bar{S}_{m,g})^2$.

The prior enters the attribution and the signal variance through offsetting factors and therefore cancels from R_g. The gain consequently depends on the labelled share, through $V_{u,g}$, and not on reading noise. The model carries the gain as the estimated moment, while the filter serves only as the D.4 diagnostic for the ordering of gains.

D.2 Measuring the Labelled Share

The labelled share of within-spell pay movement rises with income, from 0.121 at the bottom to 0.178 in the middle and 0.358 at the top. The gradient survives every specification in Table 10. I measure it on the stayer panel of Section 2, which contains 2,250,815 person-months across 103,134 persons and 127,184 spells.

For a worker i in job spell s, let y_{it} denote total earnings in month t and ℓ_{it} labelled pay, bonuses and commissions in the main specification, with spell means $\bar{y}_s$ and $\bar{\ell}_s$. I demean each within the spell and scale by spell-mean earnings, $d_{it}^{\text{tot}} = (y_{it} - \bar{y}_s)/\bar{y}_s$ and $d_{it}^{\text{lab}} = (\ell_{it} - \bar{\ell}_s)/\bar{y}_s$, and take the person-weighted coefficient from regressing the labelled deviation on the total through the origin, $\rho_g = \sum_{(i,t)\in g} w_{it}\, d_{it}^{\text{lab}}\, d_{it}^{\text{tot}} / \sum_{(i,t)\in g} w_{it}\, (d_{it}^{\text{tot}})^2 = \text{Cov}_w(d^{\text{lab}}, d^{\text{tot}})/\text{Var}_w(d^{\text{tot}})$, where w_{it} are the SIPP person weights and the sum runs over stayer person-months in group g. Seam months are dropped as in Section 2.

Writing $d^{\text{tot}} = d^{\text{lab}} + d^{\text{unlab}}$ gives $\rho_g = [\text{Var}_w(d^{\text{lab}}) + \text{Cov}_w(d^{\text{lab}}, d^{\text{unlab}})]/\text{Var}_w(d^{\text{tot}})$, which reduces to the pure variance share when the two components are uncorrelated. This appears as the alternative row of Table 10 and remains close to the main row throughout. The denominator contains permanent drift and any tenure trend, so ρ_g is a conservative proxy for the labelled share of $\sigma^2_{\theta,g}$ alone.

Specification	Bot	Mid	Top
Main: within-spell demeaned, labelled = bonus + commission	0.121	0.178	0.358
Variance share instead of regression share	0.146	0.203	0.373
Within-spell detrended	0.176	0.256	0.488
Imputed component months dropped	0.042	0.085	0.314
Overtime counted as labelled	0.134	0.197	0.370
Annual spell-year cells (different frequency)	0.050	0.060	0.075

Table 10: The measured labelled share ρ_g across specifications, by income group, SIPP stayer panel 2013–2023, person-weighted. The main row is the central specification, and the model pins its share within the band these rows span (Section 6).

The remaining rows define the measurement band. The detrended row removes a within-spell linear trend that demeaning leaves in the unlabelled residual and therefore bounds the share from above. The imputation row provides a lower bound: imputation that spreads labelled dollars across months can manufacture labelled variance, so this row drops every imputed person-month, lowering the shares while leaving the gradient intact. It is the band's lower edge because it also discards genuine labelled months. The annual row does not enter the band: at spell-year frequency, the monthly variation read by the filter has averaged out, so it measures the labelled fraction of year-to-year risk.

Receipt shares in the same records confirm the direction of the delivery gradient. Bonus receipt rises from 2.8 to 4.8 to 9.7 percent across the groups in 2018 (2.8/7.1/12.1 in 2021), while receipt of any labelled component rises from 7.9 to 16.5 percent between bottom and top. Receipt of tips falls from 2.4 to 0.9. Receipt speaks to direction rather than the level of ρ, and places the middle below the top on every labelled component.

A final check concerns the netting floor $(1 - \rho_g)\sigma^2_{\theta,g}$ that D.3 uses to score the priors. It multiplies a level-based share by the trimmed difference moment of Section 3. Measured directly as the lag-1 autocovariance of the unlabelled pay change, the floor moves by about a tenth or less under the trim in every group, while the total difference moment moves by factors of seven to fifty because the trimmed tail is almost entirely labelled bonus and commission months. The labelled share implied by the difference moments, 0.08/0.10/0.48, lies inside the band at the bottom and middle and at its detrended upper edge at the top. The top's floor comparison therefore supports the band rather than the point estimate.

D.3 The Scorecard

With the labelled share measured, the estimated belief moments of Section 5 become out-of-sample predictions. I score them using three comparisons. The floors isolate rather than predict a fourth number: the bottom's residual.

The first comparison sets the estimated priors against the netting floors $(1-\rho_g)\sigma^2_{\theta,g} = 2.21/3.63/3.74$. The middle and top priors sit on these floors: the residuals $\sigma^2_{\theta,\mathrm{perc},g} - (1-\rho_g)\sigma^2_{\theta,g} = \hat{\omega}^2_g$ of (13) are −0.68 and +0.38, with standard errors 0.94 and 0.82, respectively, both within one standard error of zero.

The middle estimate sits 0.7 standard errors below its floor, an insignificant violation of the inequality $\omega^2 \geq 0$; the model therefore carries the censored value. The bottom residual, +7.43 (standard error 1.76, $t = 4.2$), identifies the reading noise, which is concentrated where hours volatility is highest (D.5). The same comparison under the same-hours moments of Appendix C.7 gives residuals of −0.93 (0.94) at the middle, +0.22 (0.82) at the top, and +7.16 (1.76) at the bottom.

Reporting error in either survey works against the bottom's residual rather than generating it. Classical error in SIPP earnings loads one-for-one into measured $\sigma^2_{\theta,g}$ and therefore into the floors. At an error share s of transitory variance, each residual rises by $s\,\sigma^2_{\theta,g}$: the middle remains within one standard error of zero for s up to 0.37 and remains insignificant up to 0.57, while the top crosses one standard error at $s = 0.08$ and becomes significant at $s = 0.21$. The bottom's excess only grows.

On the SCE side, a common understatement of uncertainty would deflate every prior and make the excess smaller, while an overstatement large enough to manufacture the bottom residual, roughly a factor of four, would push the middle and top priors far below floors they instead sit on. The remaining group-specific overstatement is the numeracy channel conditioned on in Appendix C.4.

The second comparison uses cross-group attribution ratios rather than levels. The cell-mean shock regressor carries a worker's own shock only in part, so the slope recovers her response up to a loading that the merge itself does not pin down. This loading can be measured in SIPP: regressing each worker's own trimmed $(\Delta e_{i,t})^2$ on the cell mean of the other workers in her cell, constructed exactly as the regressor is, gives 0.43, 0.56, and 0.58 from bottom to top.[8] The loadings are similar but not identical, and the difference is conservative: correcting the ratios by them raises bottom-to-top from 3.04 to about 4.1 and moves middle-to-top from 1.76 to 1.83. Treating the loading as common therefore understates the gradient the mechanism must explain rather than

[8]The projection is the least-squares loading, which approximates the factor the system estimator recovers through lagged instruments. Its common variation is in large part the seasonal raise cycle, which under the rule is genuine updating variation. The realised squared changes are therefore responses to wage news whether or not their seasonal timing is predictable. Stripping this variation, by deseasonalising the cell series by month-of-year or by pushing the instruments beyond the smoothing window, weakens all three estimated responses and degrades the estimator's moment conditions. The news response is consequently identified largely from the seasonal arrival of wage news.

manufacturing it.

At the measured shares, the closed form (10) predicts bottom-to-top and middle-to-top attribution ratios of 1.69 and 1.31, respectively (in logs 0.52 and 0.27), against estimated ratios of 3.04 and 1.76 (in logs 1.11 and 0.57, with log standard errors 0.63 and 0.59). The gaps are $z = +0.9$ and $+0.5$, both within one standard error, although the predicted ratios are about half the estimated point gradients. Under the measured-loading correction, the estimated gradient widens and the mechanism's share falls to about a third at the bottom. I keep the uncorrected comparison as the baseline because the loading is itself an approximation, and treat the corrected comparison as the outer edge. The moment has little power against full information: the FIRE gradient $\lambda^*_{\text{Bot}}/\lambda^*_{\text{Top}} = 1.43$ (log 0.36) lies $z = 1.2$ from the estimate. The ratios therefore discipline the mechanism's direction and rough magnitude rather than establish its existence.

The third comparison is the gain diagnostic of D.4. It predicts the ordering of the gains, with the bottom slowest, but underpredicts their spread. This is based on the moments that the cadence comparison of Section 3 identifies as the least reliable. Overall, the scorecard contains two moment sets that are consistent within one standard error, one ordering that the filter gets right while underpredicting its spread, and one significant residual that the mechanism identifies rather than predicts.

D.4 The Gain Diagnostic

The filter of D.1, evaluated at the measured shares, predicts the ordering of the estimated gains but underpredicts their spread. The observation noise is $R_g = 2(\sigma^2_{\psi,g})^2(V_{u,g}/\bar{S}_{m,g})^2 =$ 262.7/229.3/168.3, which falls with the labelled share. Thus, any common drift q implies a gain that rises with income. At the common drift that best matches the three estimated gains, weighted by their precisions, $q^* = 646.5$, the predicted gains are 0.763/0.783/0.823, against estimated gains of 0.629/0.910/0.950.

The Gaussian factor of two in R_g is not load-bearing. The kurtosis of the trimmed residualised wage change is heavy but nearly common across groups (56.5/54.1/52.3), so the empirical observation noise rescales the three groups together. The refitted drift absorbs this common factor, leaving the predicted ordering and the shortfall in spread unchanged.

The diagnostic also separates constant attribution from the Bayesian rescaling of Section 4. The rescaled attributed signal has variance $2(\sigma^2_{\psi,g})^2 = 282.7/269.1/251.8$, so any common drift implies a nearly flat gain profile, at 0.79/0.79/0.80 under the same fit, against the estimated rise from 0.63 to 0.95. The constant-attribution filter instead has noise that falls with the labelled share and therefore predicts the rising order. The attribution itself cannot settle the issue, because the two rules imply coefficients on the observed squared change that differ only by the factor $V_{u,g}/\bar{S}_{m,g}$, which lies well within the standard errors (Appendix B). The evidence is therefore consistent with the constant-attribution reading rather than decisive for it: a Bayesian agent free to set her drift separately by group could match any gain profile.

Experience-based learning grades updating by age (Malmendier and Nagel, 2011), so an income gradient in the gain could instead reflect an age gradient. Within the 40-plus age band, however, the estimated gain rises from 0.57 to 0.84 to 0.98 across the income groups, and the bottom-to-top gap is over four standard errors. The under-40 band retains too few bottom-group respondents (117) to estimate the gain, so this check rests on the older band alone.

D.5 Hours by Income Group

Within-person hours volatility $\mathbb{E}[(\Delta \log h)^2]$ falls monotonically with income, at 4.10/3.06/2.36 in $(\%)^2$. This is the direction required by the reading-noise residual in D.3. The reported hours variable is heaped, so nearly all of this volatility comes from the rare schedule changes that Appendix C.7 nets out of the realised process. The gradient therefore measures exposure to schedule events and bears on the reading problem as a proxy: the within-schedule movement the worker must read through is precisely what a heaped report cannot record. Transitory shocks delivered through hours are also unlabelled, so variable hours lower ρ_g as well.

The earnings stream itself is lumpiest at the bottom. The kurtosis of the trimmed raw log earnings change falls from 46.8 through 24.1 to 15.5 across the groups, while the residualised wage change of D.4 is nearly common. The flow the worker must read is therefore more leptokurtic as well as more volatile. The check is directional only: a proportional scaling of the residual in hours volatility, calibrated at the bottom (7.43/4.10), would predict a middle residual near 5.5 against an estimated -0.68 (standard error 0.94). No intensity mapping from observables to ω^2 is therefore estimable.

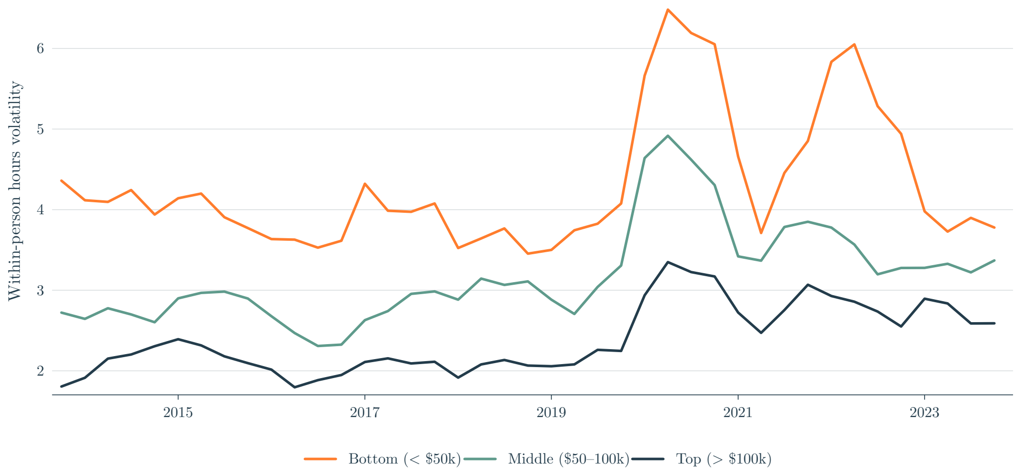

Figure 10: Within-person hours volatility $\mathbb{E}[(\Delta \log h)^2 \mid g, t]$ in $(\%)^2$ by income group. SIPP matched-stayer panel, quarterly with 4-quarter trailing MA.

E Calibrating the Belief Profile

This appendix constructs the continuous belief profile used by the model and defines the upper-bound variant reported in Section 6.

E.1 From Three Groups to a Continuum

The terciles are the finest income split the SCE can support with sufficient power, so the profiles in Figure 5 interpolate the three estimated groups rather than estimate finer bins. Re-estimating on the five native SCE brackets breaks the belief block where it is split: the lowest bracket (under \$20k, 102 respondents) has no estimable cell, the attribution turns negative in the lower-middle, the gain exceeds one in the top two brackets, and only the middle, corresponding to the unchanged tercile, survives.

Interpolation. The share ρ and gain μ are interpolated in levels, while the prior $\sigma^2_{\theta,\text{perc}}$ is interpolated in logs, using monotone cubic interpolation through the three anchors and a flat extension outside them. End tangents are zero, so each profile joins its flat extension without a kink; interior tangents are bounded so that every segment remains monotone. The prior's floor censoring at $(1 - \rho(y))\, \sigma^2_\theta(y)$ binds only at the middle anchor itself.

Table 11 evaluates the mapping at representative incomes. The attribution, discount, and permanent belief follow from the closed forms: the λ column uses the estimating form (10), while $\tilde{\lambda}$ is the implemented form based on the agent's own netted reading.

Monthly	Carried	Carried		Attribution		Discount	Perceived
income	ρ	μ	$\sigma^2_{\theta,\text{perc}}$	λ	$\tilde{\lambda}$	$V_u/\bar{S}_m$	perm. var
\$1,500	0.119	0.629	9.65	0.677	0.368	0.965	11.47
\$2,415 (Bot)	0.119	0.629	9.65	0.677	0.368	0.965	11.47
\$3,500	0.134	0.718	6.72	0.610	0.445	0.952	11.21
\$6,014 (Mid)	0.197	0.910	3.55	0.519	0.568	0.915	10.61
\$9,000	0.297	0.945	3.82	0.453	0.515	0.860	9.81
\$13,604 (Top)	0.369	0.950	4.12	0.398	0.468	0.812	9.11
\$20,000	0.369	0.950	4.12	0.398	0.468	0.812	9.11

Table 11: The income-to-belief mapping of Figure 5 evaluated at representative incomes. Rows away from the three group anchors (Bot, Mid, Top) are interpolated, not separately measured. In $(\%)^2$ except the dimensionless shares.

E.2 The Upper-Bound Variant

The belief-anchored variant takes the estimated attribution responses at face value rather than interpreting them through the measured delivery of pay. It anchors the bottom's attribution at its full-information share $\lambda^* = \sigma^2_\psi/\bar{S}_m = 0.702$ and uses the estimated attribution ratios to determine the middle and top, giving $\lambda = 0.702/0.407/0.231$. The resting permanent belief $\lambda\, \bar{S}_m$ is then 11.89/8.32/5.29, with discounts of 1.00/0.72/0.47 relative to the truth. The variant's wedge is 8.5, 19.5, and 32.4 percent of wealth and +1.2, +3.1, and +6.7 points on the marginal propensity. The shape therefore matches the measured calibration at roughly twice the size. The gap between the two calibrations captures the part of the wedge that rides on netting beyond the measured components, or on sampling noise in ratios this wide.

F Steady-State Model Output

Table 12 compares the perceived variance under the measured beliefs with FIRE by income group.

Group	$\hat{\sigma}^2_\psi$	$\sigma^2_{\theta,\text{perc}}$	Perceived $\bar{\sigma}^2_{\text{perc}}$	FIRE	Ratio
Bot	11.3	8.3	153	148	1.03
Mid	10.8	5.1	139	148	0.94
Top	9.7	3.9	124	147	0.84

Table 12: Stationary perceived variance against FIRE at the 12-month horizon, in $(\%)^2$: simulated cross-sectional means by tercile under mobility. $\hat{\sigma}^2_\psi$ is the mean permanent belief and $\sigma^2_{\theta,\text{perc}}$ the mean transitory prior, with realised counterparts 11.8/11.6/11.4 and 2.9/4.0/5.3; $\bar{\sigma}^2_{\text{perc}} = 12\hat{\sigma}^2_\psi + 2\sigma^2_{\theta,\text{perc}}$ and FIRE $= 12\sigma^2_\psi + 2\sigma^2_\theta$. Entries are computed from unrounded inputs, so displayed sums and ratios can differ in the last digit.

Table 13 gives the numbers underlying Figure 7.

	Mean wealth (months)			12-month MPC		
Group	FIRE	Belief	Δ	FIRE	Belief	Δ
Bot	2.28	2.17	−5.1%	0.470	0.477	+0.7pp
Mid	1.71	1.57	−8.3%	0.411	0.423	+1.3pp
Top	1.50	1.29	−13.8%	0.377	0.400	+2.4pp
Workers only, re-terciled among workers						
Bot	2.20	2.07	−5.8%	0.334	0.343	+0.9pp
Mid	1.54	1.39	−9.7%	0.359	0.374	+1.5pp
Top	1.47	1.25	−15.1%	0.348	0.375	+2.7pp

Table 13: Belief regime versus FIRE by income tercile. Wealth is in months of permanent income, and the 12-month MPC is the annualisation $1-(1-m)^{12}$ of the one-month response m to a transitory windfall. The regimes are paired by common random numbers. Partial equilibrium at $r = 2$ percent with common $\beta = 0.917$, which targets mean liquid wealth of about 1.8 months. The bottom tercile is 54 percent retirees, against 22 and 12 percent above; the workers-only panel re-terciles among workers. Its MPC levels sit below the full-population levels and are not monotone because removing retirees removes short horizons and high propensities, while the belief wedge itself sharpens.

Table 14 gives the general-equilibrium counterpart of Section 6.

Regime	r (% ann.)	K	H2M share
FIRE	3.00	3.65	0.338
Belief	3.21	3.45	0.342

Table 14: General-equilibrium FIRE versus belief, common $\beta = 0.917$, one-asset closure, with Z pinned from FIRE at $r^* = 3\%$ and held fixed across regimes. H2M is the hand-to-mouth share.