How do Households Respond to Perceived Income Uncertainty? Evidence from Consumption, Saving, and Debt Repayment

Xiaoting(Jasmine) Chen[1]
Leibniz Institute for Financial Research SAFE
Goethe University Frankfurt
chen@safe-frankfurt.de

Preliminary draft. Please do not circulate, distribute, or cite without permission.

August 2026

Abstract

This paper studies how households' consumption, saving, and debt repayment respond to perceived income uncertainty, separate from changes in expected income. To answer this, I develop a novel information experiment, embedded in a nationally representative panel, that holds expected income constant and manipulates only households' perceived income uncertainty. Among households already saving or carrying debt, a 20% uncertainty shock cuts consumption by 2%, raises saving by 3.9%, and increases debt repayment by 2.6%, relative to no uncertainty. A 40% shock nearly triples the consumption cut and produces correspondingly larger responses in saving and debt repayment. As the uncertainty level grows from 20% to 40%, more households adjust along all three margins. Households' balance sheet strength, measured by liquidity and debt leverage, is the main driver of this heterogeneous response. When the same households face different levels of income uncertainty captured by within-subject design, their decisions respond to the level of perceived uncertainty they currently face rather than to the path by which they arrived there. However, households respond to the size of the uncertainty shock itself, not to whether its probability was known or unknown (ambiguity).

JEL classification: D12, D81, E21, C93, D84

Keywords: Household finance, income uncertainty, subjective income distribution, precautionary saving

[1]**Acknowledgment**: I thank Christine Laudenbach and Michael Haliassos for their invaluable guidance and support throughout this project. I benefit from useful discussions with Peter Andre, Francesco D'Acunto, Olga Goldfayn-Frank, Anne-Kathrin Fischer, Tobias Schmidt, and Christian Schlag. I also thank participants at the SAFE Brownbag Seminar, 4th BCFM PhD Conference, 10th Household Finance Workshop, 2026 Frankfurt Summer School, and Bundesbank Research Center Internal Seminar. Survey support from the Deutsche Bundesbank and financial support from the SAFE Experiment Center and the Frankfurt Laboratory for Experimental Economic Research (FLEX) is gratefully acknowledged. This paper represents the author's personal views and does not necessarily reflect the views of the Deutsche Bundesbank. **Preregistration**: The study was preregistered at AEA RCT Registry https://www.socialscienceregistry.org/trials/17079

1 Introduction

> *"The* ***degree of uncertainty in future income*** *does have an important effect: consumers with greater income uncertainty,* ***ceteris paribus****, have lower current consumption."* — (Carroll, 1994)

Economic theory predicts that households with greater income uncertainty, all else equal, have lower current consumption. But future income uncertainty is not what an econometrician calibrates from realized income. It is what households themselves perceive, since they hold private information about their own income that observable characteristics cannot capture. To causally identify whether and how perceived income uncertainty itself affects current consumption is challenging as changes to perceived income uncertainty also move expected income changes in the existing panel. In addition, achieving that all-else-equal comparison requires tracing the same household's intertemporal choices as only its perceived income uncertainty changes. These three challenges leave the question open: how does perceived income uncertainty shape households' current consumption, saving, and debt repayment decisions?

I develop a novel information treatment embedded in a nationally representative survey to answer these challenges. I directly elicit households' perceived income uncertainty using probabilistic questions in an existing panel to address the first challenge. I then develop an information treatment that manipulates only households' perceived income uncertainty while keeping their expected income constant, to address the second challenge. I find that perceived income uncertainty itself shapes households' current consumption, saving, and debt repayment simultaneously. To address the third challenge — studying how the same households' intertemporal decisions respond to different uncertainty levels while holding all else unchanged — I incorporate a within-subject design into the information treatment. Households' current decisions are sensitive to changes in the uncertainty level, but what matters most is the final stage they arrive at, not the current uncertainty they experience along the way. This paper is the first information treatment that directly implements the precautionary saving theory's key comparative static — holding expected income constant while varying only its uncertainty — within the same household.

The information treatment contains four stages. At the first stage, I collect households' current consumption, saving, and debt repayment decisions before providing any information. Each allocation is reported as a share of monthly household income, with the three shares forced to sum to 100 at every stage. At the second stage, I provide every respondent with a no-uncertainty scenario, anchoring each household to its own baseline. At the third stage, I randomly assign respondents to one of two income-uncertainty sizes that manipulate only their perceived uncertainty while holding expected income fixed: a smaller, 20% shock or a larger, 40% shock. At the fourth stage, I randomly assign each respondent to one of two further treatments relative to their own stage-three scenario. The first changes only the uncertainty size, switching 20% respondents to 40% and vice versa, holding expected income constant. The second removes the known probability itself, telling respondents that the likelihood of each outcome is now unknown — an ambiguity treatment in which households rely on their own subjective beliefs.

I document four key results. The first is that *perceived income uncertainty itself matters.*

To begin, higher income uncertainty leads more households to adjust (the extensive margin). Consider a 40% uncertainty shock relative to a 20% one: higher uncertainty draws around 15% more households into adjusting in the first place. It also leads each adjusting household to adjust by more (the intensive margin). Holding expected income fixed, on average, a 20% uncertainty shock lowers the consumption share by 2.0% relative to the no-uncertainty baseline. The same uncertainty shock raises saving and debt repayment by 3.9% and 2.6%. A 40% uncertainty shock raises the consumption cut to 6.6% and produces correspondingly larger responses in saving and debt repayment, by 12.3% and 7.8% respectively. This pattern scales strongly and non-proportionally with uncertainty size: comparing the ±20% and ±40% shocks, the adjustment is disproportionately larger under the higher uncertainty level on every margin. The results also indicate that, unlike existing literature, debt repayment is a separate channel from saving: households simultaneously build liquid buffers and deleverage, a distinction panel data cannot recover since debt appears there only as a stock.

Second, when the same household faces different sizes of income uncertainty, her intertemporal decisions are sensitive to the size of that uncertainty shock. The history through which the shock arrives does not really matter. It means households' intertemporal decisions are *largely path independent*. Households re-adjust promptly when the uncertainty scenario reverses, and for most outcomes the final allocation depends primarily on the terminal uncertainty level rather than on the path through which it was reached. Across the six outcome-by-terminal-level comparisons, five show no significant history effect. The one exception is consumption after uncertainty *declines* from 40% to 20%: households that previously experienced high uncertainty remain modestly more cautious in spending than households assigned 20% uncertainty directly. This residual caution is specific to consumption and absent for saving and debt repayment.

Third, beyond these substantial average effects, the response also shows heterogeneity, driven mainly by the strength of a household's balance sheet. Under the 20% uncertainty treatment, households in the bottom quartile of cash-on-hand raise both saving and debt repayment by more than the rest of the sample. Households in the bottom quartile of the liquidity-to-income ratio, and of net wealth, do the same. Under the 40% treatment, the bottom quartile diverges further from the middle 50% and top quartile, and this pattern now holds consistently across all three liquidity variables. Leverage is a second source of heterogeneity. Households carrying secured debt (a mortgage) or consumer debt save more and pay down existing debt faster than households without such debt when they receive an uncertainty treatment. As uncertainty rises further, this saving response grows even larger, and indebted households begin cutting consumption as well.

Last, I turn to the case where the probability of income uncertainty is not known to the household. I find *little support for an ambiguity premium on average*. When the stated probabilities are withdrawn, allocations are essentially unchanged relative to the corresponding uncertainty size. Two forces jointly produce this null. The dominant pattern is that households retain their allocations: around 80% of respondents simply do not move when probabilities are withdrawn. This retention reflects path independence rather than an artifact of the design's fixed sequencing: households who see the probability structure and keep it do not respond significantly differently from households who see the probability structure and then have it withdrawn. The second is a

counter-effect among the minority who do move: their adjustments split between increases and decreases of comparable magnitude, so that offsetting

This paper mainly contributes to the precautionary-saving literature, both theoretical and empirical, as the first to use an information experiment that separates perceived income uncertainty from expected income within the same person. The theory predicts that income uncertainty raises saving through a precautionary motive (Leland, 1968; Sandmo, 1970; Carroll et al., 1992; Huggett, 1993; Carroll, 1997, 2001). But testing that prediction runs into an identification problem: uncertainty and expected income move together in observational data, so isolating uncertainty's effect requires holding expected income fixed by design. Empirically, (Carroll, 1994) supplies the first direct evidence that income uncertainty raises saving using cross-sectional data. (Carroll and Samwick, 1997) also provides direct evidence, but from panel data in which income uncertainty co-moves with other moments of the income process, so the estimated effect cannot be attributed to uncertainty alone. (Carroll and Samwick, 1998) estimates the size of the precautionary motive using only US households, leaving open whether the estimate holds elsewhere. One key feature of my vignette design isolates the second moment of income from its expected value: it directly tests whether perceived income uncertainty by itself shapes households' current decisions, the test that both the theoretical and empirical literature call for but have not yet run.

Additionally but more narrowly, this paper elicits each household's own subjective probability distribution over future income (Manski, 2004, 2018) and links those measures directly to current consumption, saving, and repayment decisions. Past empirical work instead calibrates income uncertainty from the realized variance of past income, which assumes households perceive uncertainty the way the econometrician measures it after the fact.

This design differs from what survey-based MPC studies test (Fuster and Kaplan, 2021; Bartzoka et al., 2024; Jappelli and Pistaferri, 2014; Parker and Souleles, 2019): the MPC literature asks how households would spend an additional windfall, typically motivated by fiscal policy, and does not hold expectations fixed, whereas my outcome variable directly tests households' current decisions—how they reallocate income they already have—rather than how they would spend extra income a policy might bring. The closest design is (Koşar and Melcangi, 2025), which elicits the allocation of a *windfall* rather than the reallocation of *current* income, and finds no such effect; eliciting reallocation of current income instead reveals a sharp precautionary response.

Second, this paper contributes methodologically to the literature on macro uncertainty and household decisions (Coibion et al., 2024; Weber et al., 2022; Koşar and Melcangi, 2025), which typically compares outcomes across households assigned to different scenarios. I instead use a within-subject design (Charness et al., 2012) that presents the same household with multiple scenarios, holding everything but perceived uncertainty fixed. This identifies how a given household adjusts its intertemporal decisions across scenarios, holding that household's own unobserved characteristics constant.

Third, this paper speaks to two further strands of literature. Building on this direct elicitation, I additionally test whether the mere ambiguity of an uncertainty realization matters at the household level, speaking to the macroeconomic *ambiguity-aversion* tradition (Ilut and Schneider, 2012, 2022; Ilut et al., 2025). The uncertainty-versus-ambiguity block provides household-

level evidence that, after households first observe explicitly stated probabilities, withdrawing those probabilities generates little additional adjustment on average. Rather than pricing a separate ambiguity premium, households appear to retain the subjective probabilities formed under the original uncertainty environment—though the aggregate null conceals genuine, offsetting ambiguity preferences across households. It also speaks to the *experience-effects* literature (Malmendier and Shen, 2024): the within-subject reversal delivers a causal, within-household test of whether responses depend on the history of perceived uncertainty or only its current level, and the results show behavior close to path-independent optimization, with only a narrow, consumption-specific persistence after periods of high uncertainty.

Perceived income uncertainty moves consumption, saving, and debt repayment on its own, independent of expected income, so stabilizing it is a policy margin in its own right. Because household consumption drives aggregate demand, stabilizing households' perceived income uncertainty—rather than only their expected income—would itself stabilize aggregate consumption and, through it, GDP. Communication serves this goal if it targets the uncertainty households face rather than the expected level of their income. Central banks and labor-market authorities narrow the range of income outcomes households consider plausible when they set out, before a downturn arrives, exactly how unemployment insurance, short-time work schemes, or wage subsidies will respond to one, rather than deciding support case by case after the fact. This margin matters most for households with weak balance sheets, who lack the savings buffer to self-insure and therefore cut consumption most sharply when perceived uncertainty rises.

The remainder of the paper proceeds as follows: Section 2 describes the data and vignette design, Section 3 reports the results, and Section 4 concludes.

2 Data and Design

I use the Bundesbank Online Panel–Households (BOP-HH), a monthly rotating survey of German residents aged 16 and older.[2] This nationally representative survey elicits both probabilistic income expectations and consumption expenditure from the same respondents, linking income perception directly to actual consumption. I restrict this panel to working-age individuals, aged 20 to 65, who are observed for at least nine months, since retirement shifts income and consumption patterns and within-household analysis needs repeated observations. Within this panel, I focus on respondents surveyed in November 2025 with around 3,500 participants, the wave in which a vignette experiment varying perceived income uncertainty was embedded. Table 9 reports demographics, employment, profession, income, wealth, and budget allocation for this working sample, showing that it is balanced and broadly representative of the national population.

2.1 Data Description

Eliciting subjective income distribution Do households hold genuine subjective income distributions, or do they simply pick one likely outcome? The probabilistic expectations question asks each respondent to allocate probability across income-change bins over the next 12 months.[3] Fitting a beta distribution to each respondent's full panel history shows that 87.6 percent report distributions well approximated by this flexible form, and only 2.4 percent report a point mass throughout the panel (Appendix 6.2, Figure 11): most households do hold genuine distributional beliefs. Two moments, the subjective mean and standard deviation, already summarize these beliefs, since higher-moment shape differences are small once the mean and standard deviation are held fixed (Appendix 6.2.1, Figure 12). I therefore construct perceived income growth and perceived income uncertainty as the income-normalized subjective mean and standard deviation of each respondent's distribution of expected income changes (Appendix 6.2, Equations (13), (14), (15), and (16)).

The two moments are not independent in the data. The median household perceives income uncertainty equal to 2.8 percent of current household income (Appendix 6.2, Table 7), and conditional on any given uncertainty level, roughly two-thirds of households expect income growth while the remaining third expect declines (Appendix 6.2, Figure 13; Appendix 6.2.1, Table 8). The same measured uncertainty can therefore accompany very different underlying beliefs, making it difficult to attribute consumption and saving responses to uncertainty alone using the observational panel. The resulting identification problem motivates the vignette experiment below, which varies perceived income uncertainty independently of expected income growth.

Consumption and saving variables The survey collects detailed monthly spending across the categories listed in Appendix 6.1. Total consumption is the sum of all spending categories

[2] See `https://www.bundesbank.de/en/bundesbank/research/rdsc/research-data/bop-hh-757542`.

[3] Exact wording and bins are reported in Appendix 6.1.

except saving; households allocate 86.3 percent of their monthly budget to consumption and 13.8 percent to saving (Table 9). The survey does not collect a flow measure of debt repayment, even though repayment plausibly responds to perceived income uncertainty, so the experimental design in Section 2.2 incorporates debt-repayment choices directly. The survey also collects forward-looking spending intentions, reported in Appendix 6.1 and used in later robustness checks.

Demographic and financial controls Demographic and financial controls follow standard practice in the household finance literature: age, education, household size, employment status, and home-ownership. Cash on hand follows (Jappelli and Pistaferri, 2014), the sum of liquid financial assets and current disposable income net of consumer debt. Household income and net wealth are reported in categorical brackets (Appendix 6.1); I assign each bracket's midpoint to build continuous measures, and the median household reports EUR 3,750 in monthly income and roughly EUR 131,000 in net wealth (Table 9). The two main financial controls, the liquidity-to-income ratio (liquid assets over monthly income) and the debt-to-income ratio (outstanding secured and unsecured debt over monthly income), are normalized by income rather than by wealth so they stay on the same scale as perceived income uncertainty and income growth. The median household holds 6.4 months of income in liquid assets and a debt-to-income ratio of 1.1, while the 75th percentile reaches 17 months of liquid assets and a ratio of 17.9, pointing to substantial heterogeneity in financial buffers and debt burdens (Table 9).

2.2 Information Experiment Design

This section develops a vignette experiment that identifies the causal effect of perceived income uncertainty on household consumption, saving, and debt repayment.

Two causal identification challenges Panel evidence documents a strong association between perceived uncertainty and household behavior (Appendix, Figure 13), but two problems prevent causal identification. Moment co-movement is the first: perceived income uncertainty moves together with perceived income growth, so any observed change in behavior could reflect either moment, or both. Unobserved heterogeneity is the second: panel data trace the same household over time, but unobserved factors specific to that household change together with its perceived uncertainty, breaking ceteris paribus and confounding any within-household comparison.

How the design solves both challenges Two key components of the design solve both challenges. The vignette design is pre-registered.[4] The vignette experiment is fielded to all BOP-HH panelists in November 2025. The same respondents repeatedly allocate their current monthly household income across consumption, saving, and debt repayment. The budget is normalized to 100. Each allocation follows a different income uncertainty scenario.

[4] `https://www.socialscienceregistry.org/trials/170791`

The first component is the mean-preserving spread: following (Rothschild and Stiglitz, 1970), each uncertainty scenario tells respondents their monthly household income would rise or fall by a fixed percentage with equal probability, so expected income is, by construction, identical to a guaranteed-income benchmark against which allocations are compared.

The second component is the within-subject structure: the same household is observed under more than one uncertainty magnitude, so the relevant variation is the within-person change in allocation rather than a comparison across households.

Survey flow Figure 1 summarizes the survey flow: respondents first report their current allocation with no information intervention, then complete a benchmark scenario with guaranteed income unchanged and a short uncertainty-preference module, before randomization into the four treatment arms begins.

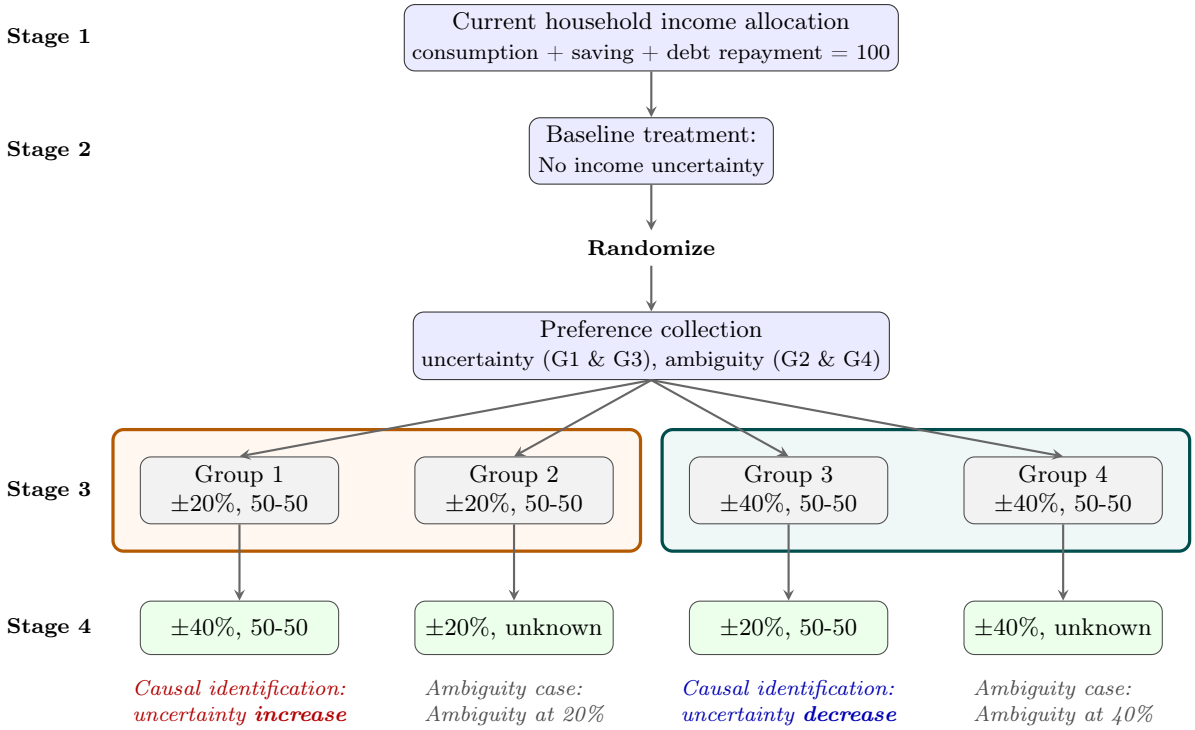


Figure 1: Survey Flow and Treatment Arms

Randomization Randomization occurs before the preference elicitation stage. Participants are first randomly assigned to one of four treatment groups, which determines the sequence of income-uncertainty scenarios they observe. Within the preference elicitation module, the presentation order of the uncertainty and uncertainty options is also randomized. This additional randomization ensures that stated preferences are not mechanically influenced by the order in which scenarios are displayed, reducing potential anchoring and order effects.

2.3 Vignette wording

Stage 1: Current household income allocation Before any information is provided, respondents report their current allocation of monthly household income across consumption,

saving, and debt repayment:

> *What share of your current monthly household income do you use for each of the following?*

Stage 2: Baseline treatment Respondents then receive a no-uncertainty scenario that serves as the baseline condition. To avoid signaling that the experiment concerns income uncertainty, the no-uncertainty scenario is framed as a guarantee rather than as zero uncertainty:

> *Please imagine that your average monthly household income over the last twelve months is* ***guaranteed to remain unchanged*** *over the next twelve months. This scenario would only apply for the next 12 months. Under these conditions, how would you* ***now*** *allocate your* ***current monthly household income****?*

The baseline scenario anchors each respondent to their own no-uncertainty allocation, preparing them for the later uncertainty treatments. Stage one and stage two questions are fielded to all panelists. Before uncertainty treatment, respondents also complete a brief preference module. Groups 1 and 3 rank income paths that differ in the magnitude of mean-preserving spreads under known probabilities; Groups 2 and 4 choose between known and unknown probabilities for a fixed income support. These measures serve as baseline characteristics for heterogeneity analysis rather than rigorous measurement. Full question wording is provided in Appendix 7.2.1.

Stage 3: isolating income uncertainty from income growth Stage 3 identifies the causal effect of perceived uncertainty on current decisions by holding expected income constant. All four arms receive a known-probability income shock: income rises or falls with equal probability, by 20 percent for Groups 1 and 2 or by 40 percent for Groups 3 and 4.

> *Please now assume that your monthly household income were to change over the next twelve months relative to your current household income as described below. This scenario would only apply for the next 12 months. In this scenario, your monthly household income would either* ***increase by 20% or decrease by 20%*** *[Groups 1 and 2] /* ***increase by 40% or decrease by 40%*** *[Groups 3 and 4]. Both scenarios would be equally likely. Under these conditions, how would you now allocate your current monthly household income?*

Stage 4: within-subject component In Stage 4, Groups 1 and 3 meet the reversed magnitude, again under explicit 50–50 probabilities. Both stages hold expected income fixed at zero, so any reallocation isolates a change in perceived income uncertainty from a change in expected income growth.

> *Please now assume that your monthly household income were to change over the next twelve months relative to your current household income as described below. This scenario would only apply for the next 12 months. In this scenario, your monthly household income would either* ***increase by 40% or decrease by 40%*** *[Group 1] /* ***increase by 20% or decrease by 20%*** *[Group 3]. Both scenarios would*

be equally likely. Under these ***new*** *conditions, how would you* ***now*** *allocate your* ***current monthly household income****?*

Group 1 moves from a 20 percent to a 40 percent uncertainty magnitude, and Group 3 moves from 40 percent down to 20 percent. Observing the same household under two different uncertainty magnitudes identifies two objects. The first is whether and how households adjust their allocation as the magnitude of income uncertainty changes, comparing the certainty-to-20-percent, certainty-to-40-percent, and reversal scenarios. The second is whether that adjustment is path-dependent — whether a household arriving at a given uncertainty level along a rising path allocates differently than one arriving along a falling path. Under a standard model, path independence is the benchmark prediction. A nonzero gap instead signals incomplete belief revision or reference dependence (Section 5.4).

For Groups 2 and 4, a further stage withdraws the probability information from the Stage 1 scenario while holding the income support fixed:

Please now assume that your monthly household income would either increase by 20% or decrease by 20% [Group 2] / increase by 40% or decrease by 40% [Group 4]. ***However, the probability of each outcome would be unknown.*** *Under these* ***new*** *conditions, how would you* ***now*** *allocate your current* ***monthly household income****?*

Since the support is unchanged, this comparison isolates the effect of unknown probabilities, or Knightian uncertainty, from the effect of uncertainty magnitude.[5] A household that carries its Stage 1 50–50 prior forward keeps its allocation unchanged; a household that instead treats the missing probabilities as ambiguous does not (Prediction 4, Section 5.4).

Decision elicitation Respondents allocate current monthly income across consumption, saving, and debt repayment, expressed as shares of a budget normalized to 100. The guaranteed-income benchmark stays visible throughout for reference (Figure 2). Holding the budget fixed across scenarios isolates reallocations from mechanical responses to different nominal income levels [6]. I elicit debt repayment separately from saving because the two reflect different balance-sheet responses. Buffer-stock models predict households build liquid buffers (saving) or pay down debt (repayment) when income is uncertain. Panel data observe debt only as a stock, so they cannot tell these responses apart.

Discussion of the design I conducted a power calculation before fielding the survey to determine the number of arms (Appendix 7.5). This calculation is also why the design stops at four arms: a fifth arm reversing the ambiguity sequence — eliciting allocations under unknown probability before revealing the known probability — would split the sample further and push its minimum detectable effect beyond a plausible economic magnitude.

[5] Knightian uncertainty refers to situations in which the probability distribution of outcomes is unknown (LeRoy and Singell, 1987).

[6] Normalizing to 100 also keeps responses comparable across households and aligns them with the income-normalized uncertainty measures used in the observational analysis (Section 2)

Question:
Please now imagine that your monthly household income (i) will change over the next twelve months relative to your current household income as described below. This scenario would only apply for the next 12 months.

In this scenario, your monthly household income (i) would either:

{if drandom2 == 1}
increase by 40% or decrease by 40%. Both scenarios would be equally likely.
{if drandom2 == 2}
increase by 20% or decrease by 20%. **However, the probability of this is unknown.**
{if drandom2 == 3}
increase by 20% or decrease by 20%. Both scenarios would be equally likely.
{if drandom2 == 4}
increase by 40% or decrease by 40%. **However, the probability of this is unknown.**

Under these new conditions, how would you now allocate your **current monthly household income (i)**?

As a reminder:
In the previous secure income scenario, you decided to spend {value of spend_guarantee_a}, save {value of spend_guarantee_b} and use {value of spend_guarantee_c} of your current monthly household income for debt repayment.

Note:
The numbers you enter for all the categories have to add up to 100.

a. For spending (i):	[Input field] percent
b. For saving (i):	[Input field] percent
c. For debt repayment (i):	[Input field] percent
Sum:	{current sum of all entered values}

-9997 Don't know
-9998 No answer

Figure 2: Example survey screen. Respondents are presented with an income-uncertainty or ambiguity scenario. They are reminded of their benchmark allocation under guaranteed income. They then allocate their current monthly household income across consumption, saving, and debt repayment. Allocations must sum to 100 percent.

3 Main Results

Perceived income uncertainty causally affects households' current financial decisions. The information treatment sample is broadly comparable to the existing panel, so the estimates that follow generalize to the population the panel describes (Section 3.1). Higher perceived uncertainty moves households along both the extensive and intensive margins, and debt repayment responds as a channel separate from saving. Only the uncertainty level a household ultimately faces matters; the path it took to reach that level does not (Section 3.3). This adjustment varies substantially across households, driven mainly by the strength of their balance sheet, in liquidity and leverage. Extending the design to ambiguity of uncertainty realization, households show a muted average response to the mere absence of known probabilities. That muted average masks heterogeneous individual responses in opposite directions that cancel out, not a genuine absence of a reaction (Section 3.4).

3.1 Sample balance and response rates

To maintain consistency between the panel analysis and the information experiment, the empirical analysis restricts attention to working-age households. Table 1 compares the full working-age BOP-HH sample with the treatment sample used in the survey experiment. It shows that the November 2025 information-treatment wave is balanced against the overall panel: once restricted to working age (column 3), the treatment sample lines up closely with the full sample on age, gender, employment status, occupation, and balance-sheet ratios, indicating that the wave in which the experiment was fielded was not itself a selected subset of the panel.

Participation in the information experiment was completely voluntary, and take-up was high: the response rate is consistent with the sample sizes assumed in the power calculations of Appendix 7.5. Attrition within the experiment is correspondingly low, and any trimming occurs almost entirely in the uncertainty ambiguity scenarios (Groups 2 and 4). Trimmed respondents are disproportionately male and without a bachelor's degree. The main analysis includes only respondents who completed all four stages.

3.2 Current and Baseline Allocation

Before any information treatment, respondents' current allocation of household income across consumption, saving, and debt repayment closely tracks their self-reported allocation in the panel, indicating that respondents have a clear sense of how they actually split their monthly spending. As shown in Figure 3, what heterogeneity that exists in the current allocation is explained mainly by socio-demographic characteristics and households' balance-sheet position, not by beliefs: neither perceived income growth nor perceived income uncertainty predicts allocation shares once these controls are included (Appendix, Table 11).

Relative to their current allocation, the no uncertainty baseline succeeds in its intended role of anchoring respondents at a common reference point of neither income uncertainty nor ambiguity. It does, however, register as a mildly negative shock on average, driven by the roughly two-thirds of respondents who expected some income growth over the next twelve months and for whom an unchanged income therefore represents a shortfall. I do not read this as undermining the

Table 1: Sample Comparison

	Total Sample	Treatment (Full)	Treatment (Work)
Demographics			
Age (mean)	48.593	55.657	48.079
Male (share)	0.635	0.605	0.631
Employment status			
Full-time	0.706	0.463	0.669
Part-time	0.172	0.120	0.165
Irregular employment	0.011	0.009	0.010
Occupation			
Salaried (baseline)	0.749	0.733	0.742
Civil servant	0.100	0.098	0.099
Self-employed	0.054	0.050	0.063
Other profession	0.079	0.053	0.072
Income and balance sheet			
Liquidity-to-income	18.297	21.223	18.210
Debt-to-income	15.771	12.279	15.101
Income uncertainty (mean)	0.050	0.045	0.051
Net wealth	317,435	375,611	349,723
Observations	44,893	3,915	2,665

Notes: Column (1) reports the full BOP-HH sample of working-age respondents (20–65) participating between February 2021 and December 2025. Column (2) restricts the sample to the November 2025 wave, when the information treatment was implemented for all panelists, and includes both working-age and retired respondents. Column (3) further restricts Column (2) to working-age respondents (20–65). Shares are reported as fractions.

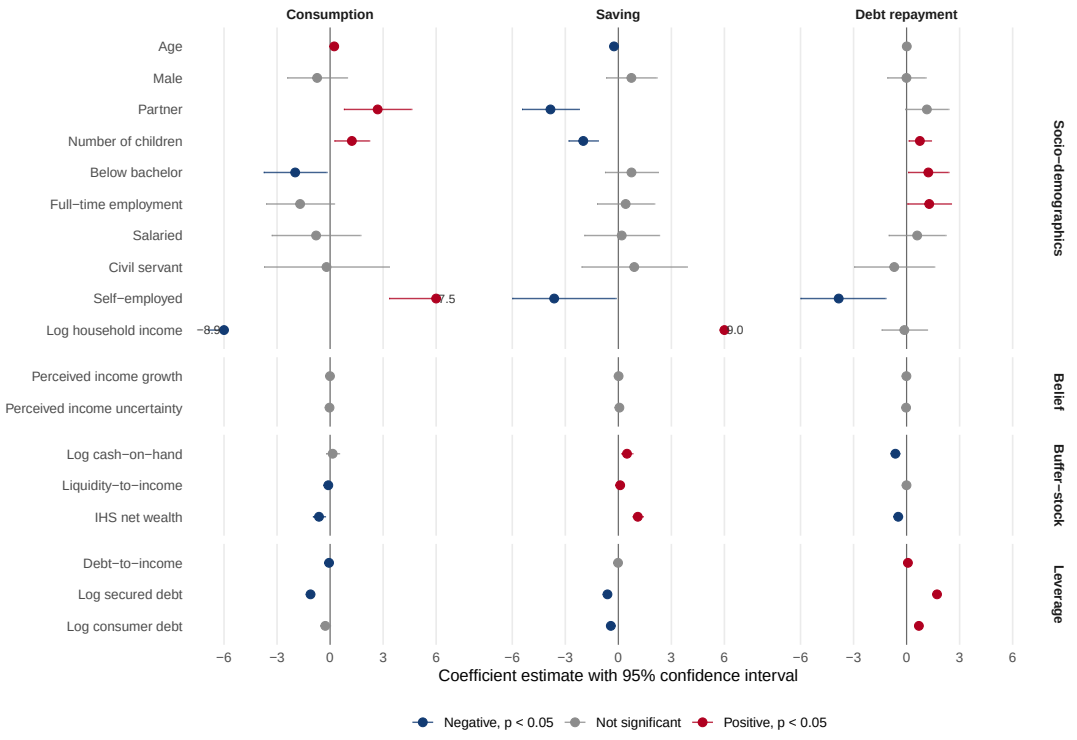


Figure 3: Heterogeneity at the current household income allocation

Notes: Full regression results (Appendix, Table 11) show that heterogeneity in the prior stage—before any information is provided—comes primarily from socio-demographics and household balance-sheet strength, measured by liquidity and leverage.

benchmark, since its purpose is only to anchor every respondent at zero uncertainty and zero ambiguity—not to match each respondent's own expected growth rate—and on that measure it succeeds. Appendix 7.4 reports the full distribution underlying the baseline result.

3.3 Perceived Income Uncertainty Matters

Average response: a clean causal estimate Table 2 reports the within-subject change in allocation moving from the certainty benchmark to a known-probability uncertainty scenario, holding expected income fixed by construction. Because the treatment is a mean-preserving spread, this comparison isolates the effect of the second moment of the subjective income distribution from the first, providing a clean causal estimate of a link that observational panel data cannot establish. Under a 20 percent uncertainty shock relative to the no-uncertainty baseline, households on average cut consumption by 2.0 percent and raise saving and debt repayment by 3.9 and 2.6 percent, respectively. Under a 40 percent shock, the same qualitative pattern holds with substantially larger magnitudes: consumption falls by 6.6 percent, saving rises by 12.3 percent, and debt repayment rises by 7.8 percent. Perceived income uncertainty alone, absent any change in expected income, is therefore sufficient to move households' current financial decisions.

Table 2: Household Responses to 20% and 40% Income Uncertainty

	Consumption	Saving	Debt Repayment
A. 20% uncertainty			
Mean log difference	-0.020***	0.039***	0.026**
	(0.005)	(0.015)	(0.011)
Paired R	0.897	0.810	0.943
Paired R^2	0.804	0.655	0.889
Observations	960	837	489
B. 40% uncertainty			
Mean log difference	-0.066***	0.123***	0.078***
	(0.008)	(0.019)	(0.015)
Paired R	0.809	0.715	0.897
Paired R^2	0.654	0.511	0.804
Observations	938	790	469

Notes: The table reports paired within-subject changes between the uncertainty and guaranteed-income scenarios. The dependent variables are log differences in consumption, saving, and debt repayment allocations. Standard errors are reported in parentheses. Paired R is the correlation between log allocations under the guaranteed-income and uncertainty scenarios; paired R^2 is its square. *, **, and *** denote significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Extensive margin: uncertainty draws more households into adjusting Perceived income uncertainty operates first on the extensive margin: raising the magnitude of uncertainty brings more households into action on each margin, not only larger adjustments among those already adjusting. Figure 4 shows that the share of households adjusting consumption rises from 46.6 to 60.9 percent as uncertainty increases from 20 to 40 percent, and the share adjusting saving rises from 48.0 to 61.4 percent. Debt repayment adjustment, among indebted households, rises from 13.9 to 20.1 percent. All three increases are statistically significant. Higher perceived

uncertainty thus recruits new households into precautionary behavior across all three margins, rather than intensifying a response that was already universal on any one of them.

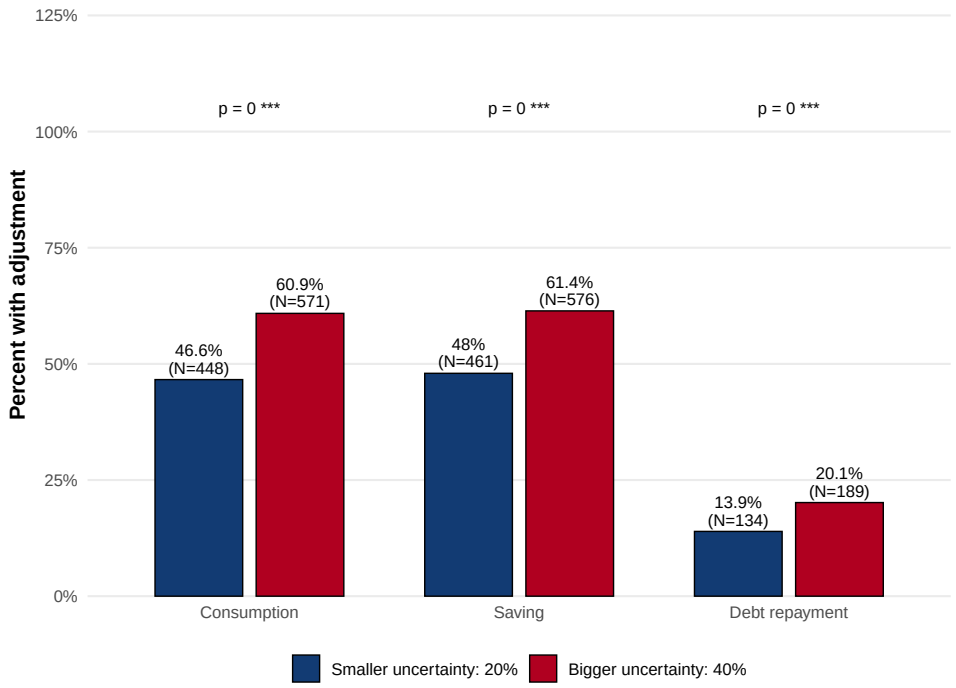


Figure 4: Extensive Margin: Share of Households Who Adjust

Notes: Bars report the share of households who adjust each allocation under the 20 percent and 40 percent uncertainty treatments. The debt-repayment sample is restricted to households with existing debt (DTI > 0). Significance stars are based on two-sample Welch t-tests comparing outcomes under 20% versus 40% uncertainty within each category.

Intensive margin: simultaneous and disproportional adjustment across all margins Perceived uncertainty also operates on the intensive margin, and does so on all three allocation margins *simultaneously* rather than one at a time: the same household that cuts consumption also raises both saving and debt repayment in response to the same uncertainty shock (Table 2). Moreover, the intensive-margin response scales more than proportionally with the size of the uncertainty. Doubling the uncertainty magnitude from 20 to 40 percent more than triples the adjustment on every margin—3.30× for consumption, 3.15× for saving, and 3.00× for debt repayment—consistent with households treating larger perceived uncertainty as calling for disproportionately stronger precautionary adjustment, in line with the prediction that the response scales with the variance of perceived uncertainty rather than its level.

Why debt repayment: two distinct channels of self-insurance That households adjust debt repayment alongside saving—rather than saving alone—indicates that precautionary behavior operates through two distinct balance-sheet channels, not one. Eliciting debt repayment as a margin separate from saving makes this visible: households build precautionary net wealth both by accumulating liquid buffers and by deleveraging, and the two channels move together rather than substituting for one another. Reducing outstanding debt lowers future liabilities and thereby serves as an additional form of self-insurance against adverse income realizations, a

channel that panel data—where debt is observed only as a stock—cannot separate from ordinary saving. Precautionary behavior in response to perceived income uncertainty is therefore best understood as an adjustment to the household's overall balance sheet, not merely as an increase in liquid saving.

Heterogeneity Having established the average response to income uncertainty, this section asks which households adjust more strongly, and why. Buffer-stock models predict that the strength of the precautionary response to uncertainty depends on a household's existing liquidity and leverage: households with little liquid buffer (Carroll, 1997; Carroll et al., 2021) or high debt have less room to absorb income shocks (Telyukova, 2013) and should therefore respond more strongly to perceived uncertainty. I focus on this balance-sheet heterogeneity because these margins exhibit clearer differences than the socio-demographic splits, which show in Figure 5.

Figure 5 reports unconditional weighted mean percent changes by balance-sheet group. The main heterogeneity comes from two household balance-sheet dimensions: leverage and buffer stocks. In the 20 percent uncertainty treatment, leverage heterogeneity is especially clear: households with higher debt-to-income reduce consumption more and increase saving and debt repayment more than households with lower debt-to-income. Buffer-stock heterogeneity is also visible, with lower cash-on-hand, liquidity-to-income, and net wealth generally associated with larger consumption reductions, although saving and debt-repayment responses are less monotone. In the 40 percent treatment, both leverage and buffer-stock differences remain present, but the ordering across groups is less systematic.

To test whether these unconditional patterns survive controlling for other household characteristics, I next estimate a conditional regression. For each outcome consumption, saving, or debt repayment (generally denoted as X) and uncertainty treatment $u \in \{20\%, 40\%\}$, I regress the log change in allocation decisions from baseline to the uncertainty scenario on socio-demographic, buffer-stock, and leverage characteristics jointly:

$$\Delta \log X_{iu} = \alpha_u + \theta_u' D_i + \beta_u' B_i + \lambda_u' L_i + \varepsilon_{iu}.$$

D_i contains socio-demographic controls (age, education, civil-servant status, self-employment status, and log household income). B_i contains buffer-stock indicators for cash-on-hand, liquidity-to-income, and net wealth, each discretized into bottom-quartile, middle-two-quartiles, and top-quartile groups, with the bottom quartile omitted. L_i contains leverage measures: debt-to-income group indicators and indicators for positive secured and consumer debt, with the no-debt group as reference. All four sets of controls enter the regression simultaneously, and I estimate this separately by outcome and treatment using weighted least squares with the Stage 1 survey weights.

The full regression results are reported in Appendix Table 18. Conditional on the full set of controls, leverage remains the strongest predictor of heterogeneity under the 20 percent treatment: relative to the no-uncertainty baseline, households in the top debt-to-income quartile cut consumption 12.3 percent more, and raise saving 18.0 percent more, than households in the bottom quartile; the middle debt-to-income group shows a similar pattern (8.8 percent more on consumption, 24.5 percent more on saving), and holding secured debt is associated with

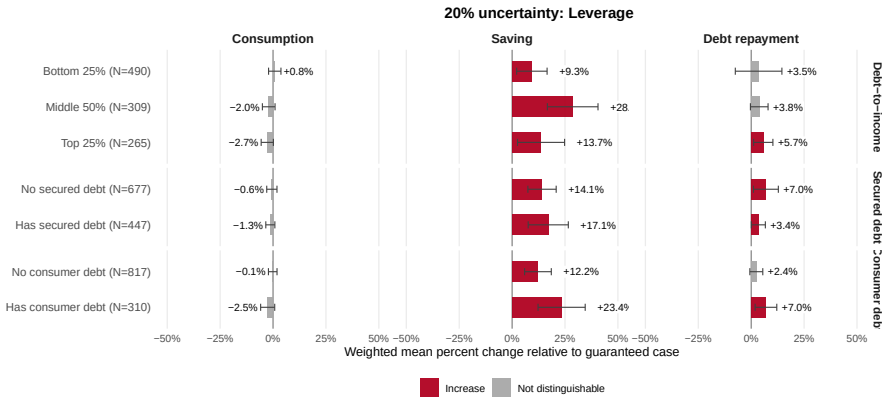
a further 7.1 percent larger consumption cut. Buffer-stock heterogeneity is present but more selective: low liquidity-to-income and low net-wealth households save significantly more, while cash-on-hand heterogeneity is only significant for saving among the lowest-holdings group. Under the 40 percent treatment, the pattern shifts: leverage effects mostly lose significance, while buffer-stock measures become more prominent, with middle net-wealth and liquidity-to-income households showing larger saving and debt-repayment responses, and having secured debt now predicting a smaller consumption cut and a larger debt-repayment increase. Socio-demographic controls remain the weakest predictors throughout. Overall, the conditional regressions confirm that leverage and buffer-stock position account for most of the heterogeneity in how strongly households respond to perceived income uncertainty, though which balance-sheet dimension dominates shifts with the size of the shock.

3.3.1 Intertemporal decisions are (mostly) path independent

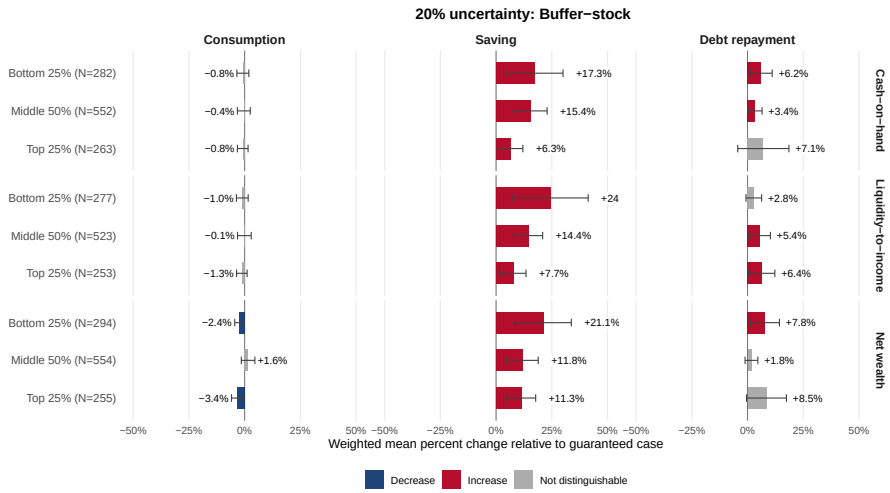
Households adjust promptly when the uncertainty scenario reverses. Households adjust their allocation promptly whenever the uncertainty scenario changes, and always in the precautionary direction. Table 3 reports the within-subject change in expenditure shares for households whose perceived uncertainty rises from 20 percent to 40 percent and, separately, for those whose uncertainty falls from 40 percent to 20 percent. When uncertainty rises, households cut consumption by 1.63 points while raising saving by 1.03 points and debt repayment by 0.61 points; when uncertainty falls, the allocation reverses, with consumption rising by 1.30 points as saving and debt repayment fall by 0.78 and 0.51 points. Every estimate is significant at the 1 percent level. The takeaway is simple: within the same household, a reversal in the uncertainty scenario produces a significant reallocation. I do not read the two legs against each other, because at this stage they begin and end at different states and are therefore not directly comparable.

Figure 6 plots the adjustments across the stages together; the response at each stage is visually clear and in line with the table. Because the second-stage legs are not directly comparable, the cleaner test of whether history matters holds the terminal uncertainty level fixed and compares households that reach it directly against those that reach it via the other level—the path-dependence test of the next subsection.

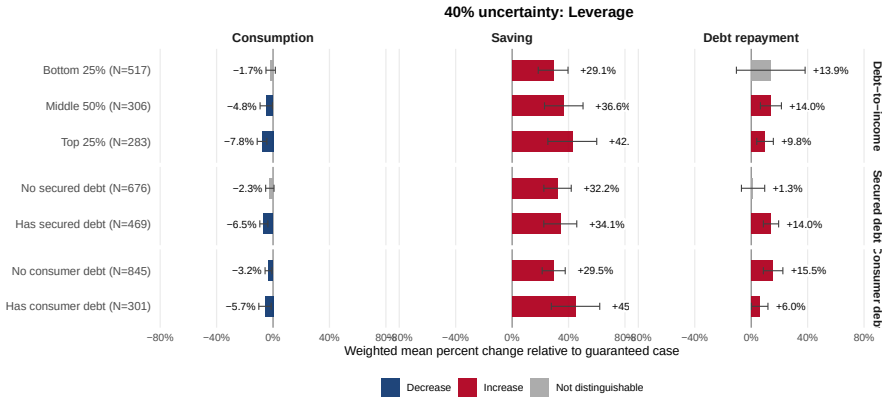
Decisions depend on the current uncertainty level, not the path. Household decisions are largely path independent: what households choose depends on the current level of perceived income uncertainty, not on the sequence through which that level is reached. Two households facing the same terminal uncertainty level should therefore choose similar consumption, saving, and debt repayment regardless of whether they arrived through increasing or decreasing uncertainty. Table 3 and Figure 7 test this by comparing, at each terminal level, households that arrive directly against those that arrive via the other level.



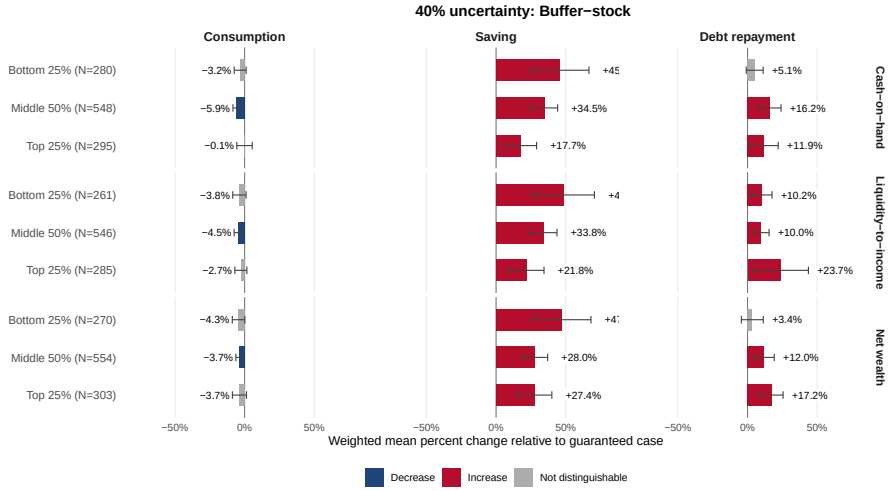
(a) Leverage: 20 percent treatment



(b) Buffer stock: 20 percent treatment



(c) Leverage: 40 percent treatment



(d) Buffer stock: 40 percent treatment

Figure 5: Balance-Sheet Heterogeneity in Uncertainty Responses

Notes: The figure reports heterogeneity in household adjustments to the 20 percent and 40 percent uncertainty treatments by balance-sheet position. Outcomes are log adjustments in consumption, saving, and debt repayment. The socio-demographic heterogeneity results are reported in the appendix. $^{+}p < 0.10$, $^{*}p < 0.05$.

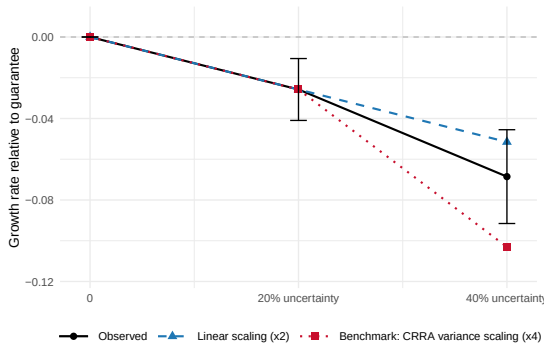
Path independence largely holds. Across the six outcome-by-terminal-level cells, five show no significant history effect. At the 40% terminal level, households arriving via 20% are statistically indistinguishable from those assigned 40% directly on every margin (differences of 0.015, −0.000, and −0.028, all insignificant); at the 20% terminal level, saving and debt repayment likewise show no significant difference. Households that reach a given uncertainty level along different paths mostly make the same decisions, as the overlapping confidence intervals in Figure 7 show. On the saving and debt margins, and at the higher terminal level on all margins, the history of perceived uncertainty does not move behaviour.

Table 3: Path Independence at Final 20 Percent and 40 Percent Uncertainty

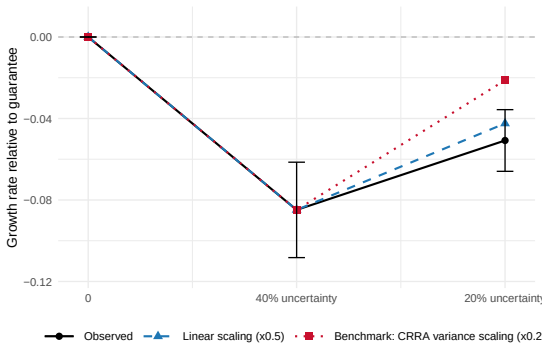
Outcome	Direct	Cumulative path	Difference
A. Final 20% uncertainty: 0–20 vs. 0–40–20			
Consumption	-0.026^{***}	-0.051^{***}	-0.025^{**}
SE	0.007	0.008	0.011
p-value	0.001	< 0.001	0.018
Obs.	474	473	–
Saving	0.076^{***}	0.096^{***}	0.019
SE	0.022	0.020	0.029
p-value	< 0.001	< 0.001	0.504
Obs.	408	424	–
Debt repayment	0.024	0.049^{***}	0.025
SE	0.015	0.018	0.023
p-value	0.104	0.005	0.277
Obs.	233	227	–
B. Final 40% uncertainty: 0–40 vs. 0–20–40			
Consumption	-0.084^{***}	-0.069^{***}	0.015
SE	0.012	0.012	0.017
p-value	< 0.001	< 0.001	0.357
Obs.	479	446	–
Saving	0.151^{***}	0.151^{***}	−0.000
SE	0.027	0.031	0.041
p-value	< 0.001	< 0.001	0.998
Obs.	407	363	–
Debt repayment	0.074^{***}	0.046^{*}	−0.028
SE	0.021	0.024	0.032
p-value	< 0.001	0.060	0.382
Obs.	233	219	–

Notes: Each panel compares household responses conditional on reaching the same terminal uncertainty level through different histories. The "Difference" column reports the cumulative-path estimate minus the direct-path estimate. The reported p-values correspond to the direct-path estimate, cumulative-path estimate, and path-dependence difference, respectively. SE denotes standard error. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

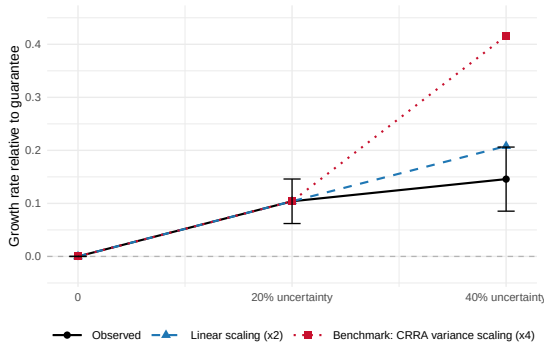
One exception: residual caution in consumption. The main exception appears in consumption at the 20% terminal uncertainty level (Figure 7a). Households that previously experienced 40% uncertainty reduce consumption by an additional 2.5 log points relative to households directly assigned 20% uncertainty (−0.051 versus −0.026; difference −0.025, $p < 0.05$). Precautionary restraint accumulated under high uncertainty does not fully unwind once uncertainty recedes. Importantly, the same pattern does not appear for saving or debt repayment, where the history terms are small and statistically insignificant.



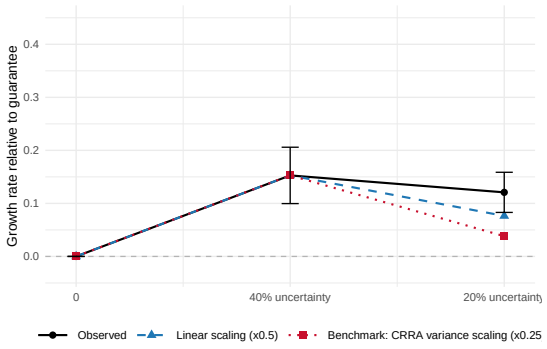
(a) Consumption: Increasing uncertainty



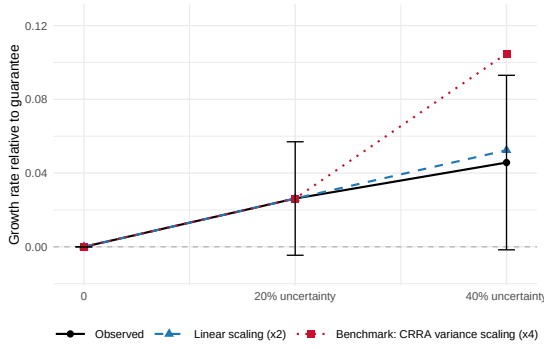
(b) Consumption: Decreasing uncertainty



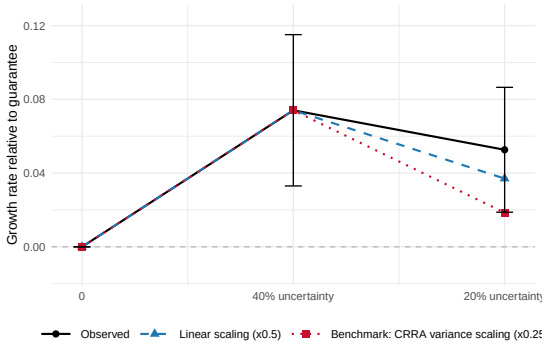
(c) Saving: Increasing uncertainty



(d) Saving: Decreasing uncertainty



(e) Debt Repayment: Increasing uncertainty

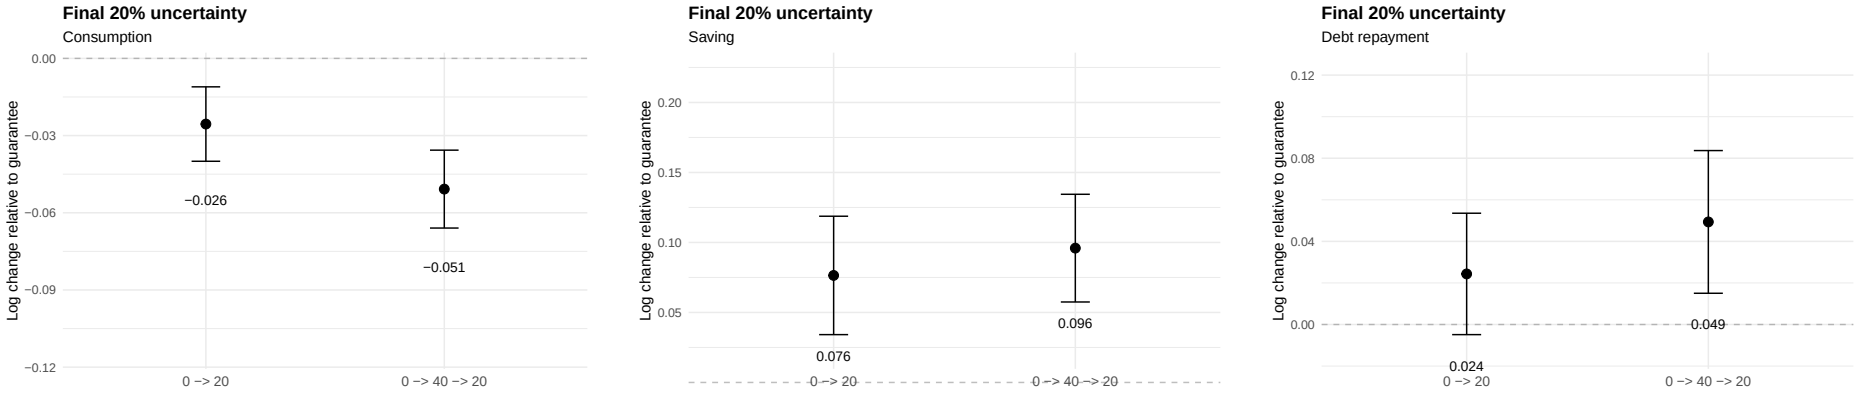


(f) Debt Repayment: Decreasing uncertainty

Figure 6: Adjustment to Different Income-uncertainty Scenarios

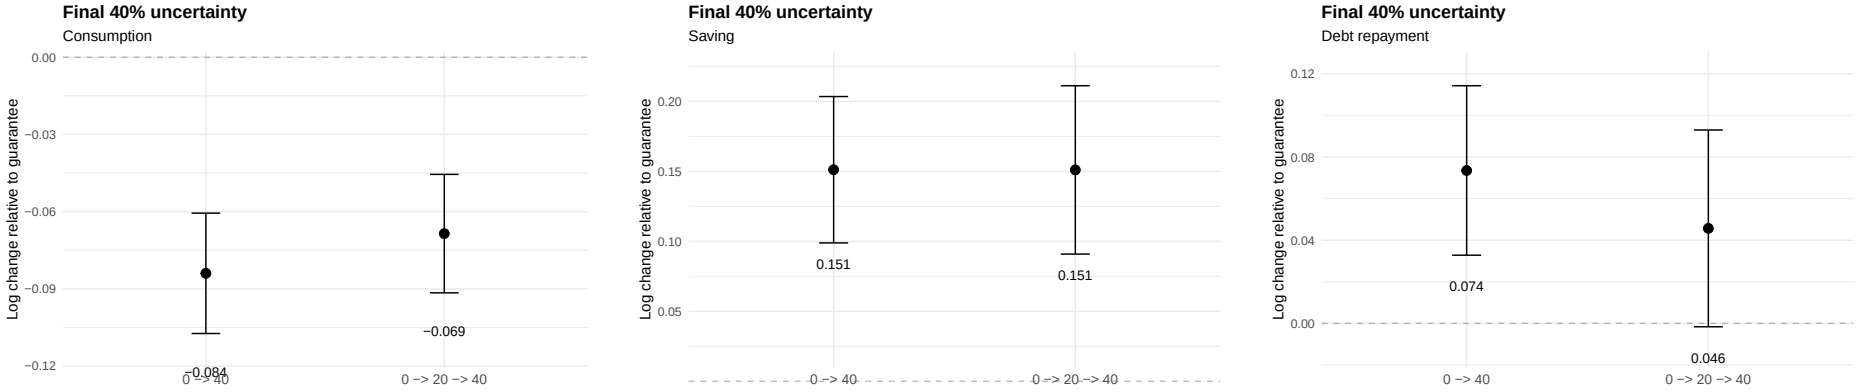
Notes: The figures first show that households are sensitive to the size of perceived income uncertainty: responses to the 40 percent treatment are systematically larger than responses to the 20 percent treatment. The figures then compare these observed responses to benchmark scaling predictions under a standard CRRA precautionary saving model, with the derivation detailed in Appendix 5. Panels on the left show responses when perceived income uncertainty increases from 20 percent to 40 percent, while panels on the right show responses when uncertainty decreases from 40 percent to 20 percent. The "Observed" line reports empirical responses, the "Linear scaling" line shows proportional scaling with the change in perceived uncertainty, and the "Benchmark: CRRA variance scaling" line shows the theoretical prediction under variance-based precautionary motives. Across all outcomes, observed responses are substantially smaller than variance scaling predictions and closer to linear scaling.

A. Final 20 percent Perceived Income uncertainty



(a) Consumption (b) Saving (c) Debt repayment

B. Final 40 percent Perceived Income uncertainty



(d) Consumption (e) Saving (f) Debt repayment

Figure 7: Path Independence under Alternative uncertainty Histories

Notes: The figure compares cumulative household responses ending at the same perceived income uncertainty level under different uncertainty histories. Panel A considers the final 20 percent uncertainty, while Panel B considers the final 40 percent uncertainty. In each panel, the left bar reports the direct response to the final uncertainty level, whereas the right bar reports the cumulative response when households first experience a different uncertainty level before transitioning to the same final state. Despite facing identical final uncertainty levels, households exposed to different prior uncertainty histories do not exhibit systematically different responses, indicating path independence in household financial behavior.

3.3.2 Robustness

I subject both the stage-one (level) responses and the stage-two (sequential) responses to a battery of robustness checks; full procedures and tables are in Appendices 8 and 9.

Stage-three responses. Three concerns could undermine the baseline estimates. First, the within-household responses contain many zeros, skewness, and outliers, so I re-test each outcome with a Wilcoxon signed-rank test; all three outcomes reject the null of no effect at conventional levels for both shocks ($p < 0.001$ for the 40 percent shock, $p \leq 0.002$ throughout), confirming the results without distributional assumptions. Second, to rule out that the effects reflect treatment-group composition, I regress each outcome on demographic and financial controls; almost no subgroup coefficient is significant and the few that are show no systematic pattern, so the effects are not an artifact of sample composition. Third, because responses are measured relative to each household's certainty allocation, baseline allocations could drive them mechanically; although statistically associated with the responses, they explain less than 4 percent of the variation ($R^2 < 0.04$), ruling out mechanical rescaling.

Sequential (stage-four) responses. The sequential adjustments survive the same scrutiny. A Wilcoxon signed-rank test confirms that both the upward ($20 \rightarrow 40$) and downward ($40 \rightarrow 20$) adjustments are significant across all three outcomes ($p < 0.001$ to 0.035). An order-effect test comparing groups that reach the same uncertainty level in different presentation orders finds no significant difference for any outcome (all but one marginal case at $p > 0.05$), ruling out order-driven bias. Among households that do adjust, a sign test shows the direction of change is overwhelmingly consistent with the shock—for example, 59 percent of savers increase saving when uncertainty rises and 67 percent of spenders cut consumption when it falls. Finally, a magnitude (ratio) test confirms the adjustments are economically meaningful: saving rises by 17.8 percent and debt repayment by 29.2 percent relative to the initial allocation in the upward group, with signs reversing symmetrically in the downward group.

3.4 Ambiguity of uncertainty realization

This subsection asks whether households respond to Knightian uncertainty as they respond to known uncertainty. On average, they do not: withdrawing explicit probabilities while holding the income support fixed generates essentially no change in allocation, a sharp contrast to the large and significant responses to known uncertainty documented above. This null average, however, does not indicate that ambiguity is behaviorally irrelevant. Most households retain their existing allocation when probabilities are withdrawn, a form of inertia that proves genuine rather than an artifact of the design's fixed sequencing; among the minority who do adjust, responses divide between increases and decreases of comparable magnitude, so offsetting behavior accounts for much of the null average. The determinants of who adjusts, and in which direction, are tied less to stated preferences over ambiguity than to the same perceived-uncertainty and balance-sheet channels that amplify the response to known uncertainty.

3.4.1 Ambiguity barely moves the average

Table 4 traces the full progression: first from no uncertainty to a known income-uncertainty probability, then from known to unknown probabilities, and finally the cumulative response across the sequence. The first stage shows strong precautionary responses once households are presented with explicit probabilities of income-uncertainty realization. Moving from guaranteed income to a known 20 percent uncertainty environment reduces consumption and raises debt repayment, while the known 40 percent treatment produces substantially larger adjustments across all three outcomes: consumption falls by 4.7 percent, saving increases by 9.4 percent, and debt repayment rises by 8.3 percent. Once households perceive explicit income uncertainty, they immediately strengthen their balance sheets.

When the same uncertainty level is subsequently presented *without* explicit probabilities, however, households barely adjust. The estimated effects of removing probabilities are economically small and statistically insignificant across all three outcomes at both the 20 percent and 40 percent levels; consumption, saving, and debt repayment remain essentially unchanged when households move from known to unknown probabilities.

The cumulative response reinforces this: the total response after moving from guaranteed income to known uncertainty and then to unknown probabilities is nearly identical to the response generated by the known-uncertainty treatment alone. The economically meaningful adjustment occurs when households first learn the probability structure of income uncertainty, not when probabilities later become unspecified.

3.4.2 Why is the average response to ambiguity so muted?

The near-zero average response documented above admits two explanations that are not mutually exclusive: either most households simply do not revise their allocation when probabilities are withdrawn. Or households do revise, but in offsetting directions that cancel in the aggregate. I examine both.

Most households do not move. Figure 8 shows the distribution of responses when probabilities move from known to unknown. Across all six outcome-uncertainty combinations, a large majority of respondents—between 75 and 87 percent—report no change in allocation at all; only 8 to 14 percent increase and a similar share decrease. This pattern could in principle reflect the specific order in which my design withdraws probabilities: because statistical power required fixing the sequence as known-then-unknown rather than randomizing the order, a respondent who simply anchors on the first-seen 50–50 split would appear not to move for mechanical reasons. We check this using the between-group comparison, which shows that households who never see an explicit probability respond to ambiguity no differently on average than those who do; inertia is therefore a genuine feature of the response to ambiguity, not an artifact of moving from a 50–50 anchor.

Among those who do move, responses offset. Figure 9 decomposes the average log change into the contribution of decreasers and increasers separately. At every outcome-uncertainty combination, the two contributions are similar in magnitude and opposite in sign—for example,

Table 4: From Known to Unknown: Persistent Subjective Probabilities

	Consumption	Saving	Debt repayment
Stage 3: Guaranteed → known uncertainty			
Known prob. 20%	-0.014^{*}	0.003	0.028^{*}
	(0.008)	(0.021)	(0.015)
R_t^2	0.007	0.000	0.014
Observations	486	429	256
Known prob. 40%	-0.047^{***}	0.094^{***}	0.083^{***}
	(0.010)	(0.027)	(0.021)
R_t^2	0.044	0.031	0.064
Observations	459	383	236
Stage 4: Known → unknown uncertainty			
Unknown prob. 20%	−0.004	0.017	−0.006
	(0.006)	(0.014)	(0.012)
R_t^2	0.001	0.004	0.001
Observations	452	402	239
Unknown prob. 40%	−0.001	−0.005	0.006
	(0.007)	(0.018)	(0.017)
R_t^2	0.000	0.000	0.001
Observations	418	347	212
Cumulative response: Guaranteed → unknown uncertainty			
$0 \rightarrow 20 \rightarrow$ unknown	-0.019^{**}	0.032	0.021
	(0.008)	(0.020)	(0.014)
R_t^2	0.012	0.007	0.010
Observations	451	400	238
$0 \rightarrow 40 \rightarrow$ unknown	-0.048^{***}	0.104^{***}	0.079^{***}
	(0.010)	(0.024)	(0.017)
R_t^2	0.054	0.052	0.091
Observations	418	353	215

Notes: Notes: The table reports within-subject log changes in household financial allocations across sequential uncertainty treatments. Stage 1 compares guaranteed income with known-uncertainty environments with explicit probabilities. Stage 2 measures the additional response when probabilities are removed while holding the income-uncertainty environment fixed. The final panel reports cumulative responses from guaranteed income to ambiguity. Under ambiguity aversion, households should exhibit additional precautionary adjustment once probabilities become unknown; the small and insignificant Stage 2 estimates provide little evidence of such an ambiguity premium. $R_t^2 = t^2/(t^2 + df)$ is reported as a t-test effect-size measure, not regression fit. Stage 1 uses only the ambiguity arm, Groups 2 and 4, whereas Table **??** pools both arms at each uncertainty level. Standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

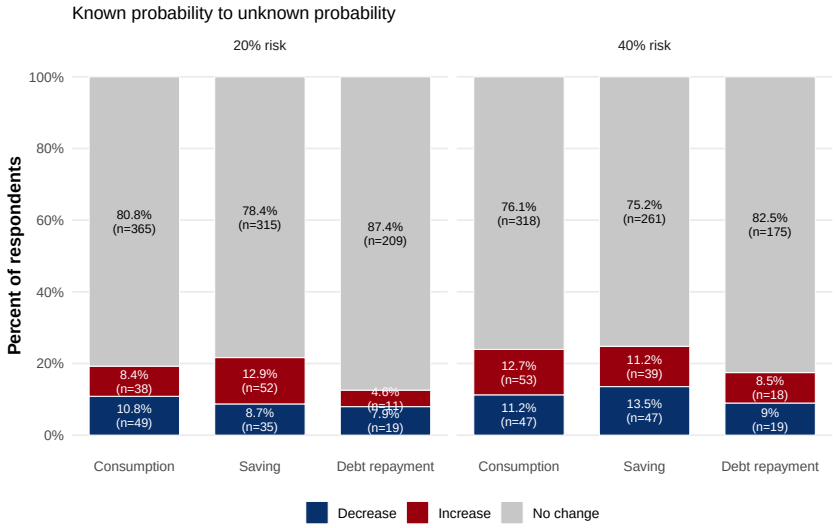


Figure 8: Adjustment Direction under Ambiguity

Notes: Bars show the share of respondents who decrease, increase, or leave unchanged their allocation when probabilities move from known to unknown, by outcome and uncertainty size.

at 40 percent uncertainty, consumption decreases contribute −0.069 and increases contribute +0.064 to the mean change. The muted average therefore is not only a story of inertia: a non-trivial minority of households do adjust substantially in response to ambiguity, and their adjustments largely cancel out rather than reinforcing each other.

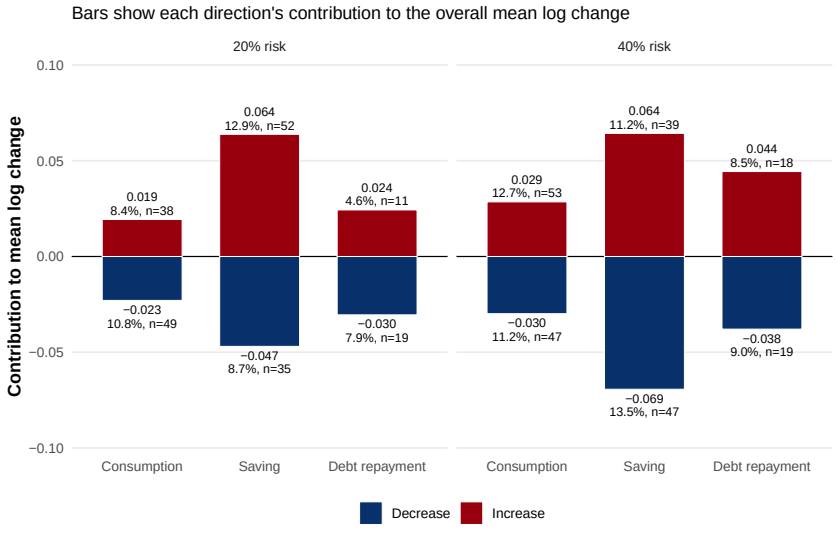


Figure 9: Contribution of Decreasers and Increasers to the Mean Log Change

Notes: Bars show each direction's contribution to the overall mean log change in allocation when probabilities move from known to unknown, by outcome and uncertainty size.

3.4.3 What explains the counter-effects?

Then what distinguishes households who decrease from those who increase? Ambiguity preference—elicited directly, and used in Section 3.3 for other purposes—explains this only selectively. The stronger and more consistent predictors are perceived income uncertainty and balance-sheet

position, the same belief and buffer-stock channels identified in the uncertainty analysis above.

For each outcome and uncertainty level, I estimate

$$Y_i = \alpha + \pi' \text{AmbigPref}_i + \beta \, \text{ExpChange}_i + \rho \, \text{Incomeuncertainty}_i + \gamma' \text{Demographics}_i + \delta' \text{BalanceSheet}_i + \varepsilon_i,$$

where Y_i is an indicator for the direction of the second-stage response (any adjustment, decrease, or increase) or the absolute log change among adjusters; AmbigPref_i is the respondent's stated preference for known versus unknown probabilities; ExpChange_i and $\text{Incomeuncertainty}_i$ are the respondent's expected income change and perceived income uncertainty prior to treatment; Demographics_i includes age, gender, education, and employment status; and BalanceSheet_i includes log household income, log cash-on-hand, inverse-hyperbolic-sine net wealth, liquidity-to-income, and debt-to-income. Table 5 reports the coefficients on $\text{Incomeuncertainty}_i$, BalanceSheet_i, and AmbigPref_i from this regression.

Table 5: Predictors of Adjustment Direction and Magnitude under Ambiguity

Spread level	Outcome	Predictor	Direction outcome	Coefficient
20%	Consumption	Perceived income uncertainty	Pr(any adjustment)	0.730^{**}
		Perceived income uncertainty	Pr(decrease)	0.547^{**}
		Perceived income uncertainty	Pr(increase)	0.183
20%	Saving	Log household income	Pr(any adjustment)	0.071^{*}
		Log cash-on-hand	Magnitude (adjusters)	-0.126^{**}
40%	Consumption	Prefer unknown prob.	Magnitude (adjusters)	-0.094^{*}
		Log cash-on-hand	Magnitude (adjusters)	0.065^{**}
		IHS net wealth	Magnitude (adjusters)	-0.028^{***}
		Debt-to-income	Magnitude (adjusters)	0.002^{**}
40%	Saving	Prefer unknown prob.	Magnitude (adjusters)	0.231^{**}
40%	Debt repayment	IHS net wealth	Pr(any adjustment)	-0.020^{***}
		IHS net wealth	Pr(increase)	-0.016^{***}

Notes: Coefficients are from linear probability (direction outcomes) and OLS (magnitude among adjusters) models of the second-stage response to ambiguity, $\Delta_{\text{uncertainty},i} = \log(\text{unknown}_i) - \log(\text{known}_i)$, controlling for demographics and balance-sheet variables not shown; see Table 23 for the full specification. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

At 20 percent ambiguity, higher perceived income uncertainty predicts a higher probability of cutting consumption specifically (not of adjusting in either direction), a pattern consistent with genuine precautionary behavior rather than noise. At 40 percent ambiguity, balance-sheet position becomes more prominent: cash-on-hand, net wealth, and the debt-to-income ratio all predict the *magnitude* of adjustment among consumption and debt-repayment adjusters, and stated ambiguity preference predicts adjustment magnitude for saving specifically. Overall, the evidence points to a rough hierarchy—perceived income uncertainty first, balance-sheet position second, ambiguity preference selectively, and demographics last—that mirrors the belief-channel result from the uncertainty analysis: the same underlying beliefs that amplify the response to known uncertainty also govern who moves, and in which direction, once probabilities become ambiguous.

Appendix 10 reports the full regression results underlying this ranking.

4 Conclusion and Discussion

This paper asks how households adjust current consumption, saving, and debt repayment when perceived income uncertainty changes, holding expected income fixed. I answer this question with a novel vignette experiment embedded in a nationally representative German household survey of roughly 2,500 working-age respondents. Three challenges have stood in the way of a causal answer. First, the future income households perceive is not the same as the one econometricians calibrate. Second, in existing panels that elicit perceived income uncertainty, it moves together with expected income. Third, unobservable heterogeneity makes it hard to trace the same households' intertemporal choices while holding all other conditions unchanged.

I address these challenges with a within-subject randomized vignette experiment. Mean-preserving spreads vary perceived uncertainty while holding expected income fixed, breaking the first-moment/second-moment co-movement at its source. The within-subject design lets each household serve as its own control. This ensures that all other conditions remain unchanged, focusing only on how the same household's intertemporal decisions respond to different income uncertainty scenarios. A further case study withdraws the probability of the uncertainty realization to test how ambiguity about that realization (Knightian uncertainty) influences households' decisions.

Three results follow. First, perceived income uncertainty alone causes households to cut consumption, build saving, and accelerate debt repayment, on both the extensive and intensive margins: a 40 percent shock draws 15 percent more households into adjusting than a 20 percent shock, and among those who adjust, the response roughly triples in size as the shock doubles. This heterogeneity in response traces mainly to a household's balance sheet, its liquidity and leverage. Second, the same households adjust to the size of uncertainty they face but not to how they arrived at it: the final allocation depends on the terminal uncertainty level, not the path to it, with only a modest residual caution in consumption. Third, households do not, on average, distinguish uncertainty from Knightian uncertainty: withdrawing stated probabilities leaves allocations essentially unchanged, because genuine inertia among most households offsets active adjustment among the ambiguity-averse minority who do react.

These findings give causal evidence that subjective income uncertainty reshapes household financial decisions independently of expected income, in proportion to its variance rather than its level, and in proportion to a household's own financial buffers. The result also carries a policy implication: because perceived uncertainty matters for consumption in its own right, stabilizing household spending requires stabilizing perceived income uncertainty, not only expected income; communication that credibly compresses that uncertainty could itself temper precautionary cuts in consumption.

References

Louiza Bartzoka, Olga Goldfayn-Frank, and Nate Vellekoop. Gender Differences in Reactions to Income Shocks, November 2024. URL https://papers.ssrn.com/abstract=5042626.

C. D. Carroll. How does Future Income Affect Current Consumption? *The Quarterly Journal of Economics*, 109(1):111–147, February 1994. ISSN 0033-5533, 1531-4650. doi: 10.2307/2118430. URL https://academic.oup.com/qje/article-lookup/doi/10.2307/2118430.

Christopher D. Carroll. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics*, 112(1):1–55, 1997. ISSN 0033-5533. URL https://www.jstor.org/stable/2951275.

Christopher D. Carroll. Precautionary Saving and the Marginal Propensity to Consume out of Permanent Income, April 2001. URL https://www.nber.org/papers/w8233.

Christopher D. Carroll and Andrew A. Samwick. The nature of precautionary wealth. *Journal of Monetary Economics*, 40(1):41–71, September 1997. ISSN 0304-3932. doi: 10.1016/S0304-3932(97)00036-6. URL https://www.sciencedirect.com/science/article/pii/S0304393297000366.

Christopher D Carroll and Andrew A Samwick. How Important is Precautionary Saving? *The Review of Economics and Statistics*, 1998.

Christopher D. Carroll, Robert E. Hall, and Stephen P. Zeldes. The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, 1992(2): 61, 1992. ISSN 00072303. doi: 10.2307/2534582. URL https://www.jstor.org/stable/2534582?origin=crossref.

Christopher D. Carroll, Martin B. Holm, and Miles S. Kimball. Liquidity constraints and precautionary saving. *Journal of Economic Theory*, 195:105276, July 2021. ISSN 0022-0531. doi: 10.1016/j.jet.2021.105276. URL https://www.sciencedirect.com/science/article/pii/S0022053121000934.

Gary Charness, Uri Gneezy, and Michael A. Kuhn. Experimental methods: Between-subject and within-subject design. *Journal of Economic Behavior & Organization*, 81(1):1–8, January 2012. ISSN 0167-2681. doi: 10.1016/j.jebo.2011.08.009. URL https://www.sciencedirect.com/science/article/pii/S0167268111002289.

Olivier Coibion, Dimitris Georgarakos, Yuriy Gorodnichenko, Geoff Kenny, and Michael Weber. The Effect of Macroeconomic Uncertainty on Household Spending. *American Economic Review*, 114(3):645–677, March 2024. ISSN 0002-8282. doi: 10.1257/aer.20221167. URL https://www.aeaweb.org/articles?id=10.1257/aer.20221167.

Andreas Fuster and Greg Kaplan. What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans. *Review of Economic Studies*, 0:1–36, 2021.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, January 1989. ISSN 0304-4068. doi: 10.1016/0304-4068(89)90018-9. URL https://www.sciencedirect.com/science/article/pii/0304406889900189.

Luigi Guiso, Tullio Jappelli, and Terlizzese Daniele. Earnings uncertainty and precautionary saving. *Journal of Monetary Economics*, 30(2):307–337, November 1992.

Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5):953–969, September 1993. ISSN 0165-1889. doi: 10.1016/0165-1889(93)90024-M. URL https://www.sciencedirect.com/science/article/pii/016518899390024M.

Cosmin Ilut and Martin Schneider. Ambiguous Business Cycles. *NBER Working Paper*, 2012.

Cosmin Ilut and Rosen Valchev. Economic Agents as Imperfect Problem Solvers. *The Quarterly Journal of Economics*, 138(1):313–362, February 2023. ISSN 0033-5533. doi: 10.1093/qje/qjac027. URL https://doi.org/10.1093/qje/qjac027.

Cosmin L Ilut and Martin Schneider. Modeling Uncertainty as Ambiguity: a Review. 2022.

Cosmin L Ilut, Ralph Luetticke, and Martin Schneider. HANK's Response to Aggregate Uncertainty in an Estimated Business Cycle Model. *NBER Working Paper*, No. 33331, 2025.

Tullio Jappelli and Luigi Pistaferri. Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–136, October 2014. ISSN 1945-7707. doi: 10.1257/mac.6.4.107. URL https://www.aeaweb.org/articles?id=10.1257/mac.6.4.107.

Miles S. Kimball. Precautionary Saving and the Marginal Propensity to Consume, July 1990. URL https://www.nber.org/papers/w3403.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A Smooth Model of Decision Making under Ambiguity. *Econometrica*, 73(6):1849–1892, 2005. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2005.00640.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2005.00640.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2005.00640.x.

Gizem Koşar and Davide Melcangi. Subjective Uncertainty and the Marginal Propensity to Consume, March 2025. URL https://papers.ssrn.com/abstract=5220724.

Hayne E. Leland. Saving and Uncertainty: The Precautionary Demand for Saving. *The Quarterly Journal of Economics*, 82(3):465–473, 1968. ISSN 0033-5533. doi: 10.2307/1879518. URL https://www.jstor.org/stable/1879518.

Stephen F. LeRoy and Larry D. Singell. Knight on Risk and Uncertainty. *Journal of Political Economy*, 95(2):394–406, 1987. ISSN 0022-3808. URL https://www.jstor.org/stable/1832078.

Ulrike Malmendier and Leslie Sheng Shen. Scarred Consumption. *American Economic Journal: Macroeconomics*, 16(1):322–355, January 2024. ISSN 1945-7707, 1945-7715. doi: 10.1257/mac.20210387. URL https://pubs.aeaweb.org/doi/10.1257/mac.20210387.

Charles F. Manski. Measuring Expectations. *Econometrica*, 72(5):1329–1376, September 2004. ISSN 0012-9682, 1468-0262. doi: 10.1111/j.1468-0262.2004.00537.x. URL http://doi.wiley.com/10.1111/j.1468-0262.2004.00537.x.

Charles F. Manski. Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise. *NBER Macroeconomics Annual*, 32:411–471, April 2018. ISSN 0889-3365. doi: 10.1086/696061. URL https://www.journals.uchicago.edu/doi/10.1086/696061.

Jonathan A. Parker and Nicholas S. Souleles. Reported Effects versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates. *American Economic Review: Insights*, 1(3):273–290, December 2019. ISSN 2640-205X. doi: 10.1257/aeri.20180333. URL https://www.aeaweb.org/articles?id=10.1257/aeri.20180333.

Michael Rothschild and Joseph E Stiglitz. Increasing risk: I. A definition. *Journal of Economic Theory*, 2(3):225–243, September 1970. ISSN 0022-0531. doi: 10.1016/0022-0531(70)90038-4. URL https://www.sciencedirect.com/science/article/pii/0022053170900384.

A. Sandmo. The Effect of Uncertainty on Saving Decisions. *The Review of Economic Studies*, 37(3):353–360, 1970. ISSN 0034-6527. doi: 10.2307/2296725. URL https://www.jstor.org/stable/2296725.

Irina A. Telyukova. Household Need for Liquidity and the Credit Card Debt Puzzle. *The Review of Economic Studies*, 80(3):1148–1177, July 2013. ISSN 1467-937X, 0034-6527. doi: 10.1093/restud/rdt001. URL https://academic.oup.com/restud/article-lookup/doi/10.1093/restud/rdt001.

Michael Weber, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion. The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications. *Journal of Economic Perspectives*, 36(3):157–184, August 2022. ISSN 0895-3309. doi: 10.1257/jep.36.3.157. URL https://www.aeaweb.org/articles?id=10.1257/jep.36.3.157.

5 Appendix: Stylized model

This section provides a stylized two-period model of household budget allocation under exogenous variation in income. Its role is interpretive: each vignette (Section 2.2) is mapped to an exogenous shift in subjective beliefs, and the model translates that shift into testable predictions for consumption, saving, and debt repayment. Derivations are collected in Appendix 5.6.

5.1 Environment and preferences

A household lives for two periods. In period 1 ("now"), it enters with liquid assets a_0, debt b_0, and income y_0, and chooses consumption c, end-of-period assets $a' \geq 0$, and debt repayment $x \in [0, Rb_0]$ subject to

$$c + a' + x = y_0 + R\, a_0, \tag{1}$$

where $R = 1 + r$ is the common gross return on assets and debt.[7] In period 2 the household draws stochastic income y' from a subjective distribution F and consumes its remaining resources. Non-consumption resources split between saving s and repayment d; under $R_a = R_b$ the two are perfect substitutes for the aggregate prediction, and I treat them separately only for description.

Preferences follows CRRA utility with discount factor $\beta \in (0, 1)$:

$$U = u(c) + \beta\, \mathbb{E}_F\big[u(c')\big], \qquad u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad \gamma > 0. \tag{2}$$

CRRA implies prudence ($u''' > 0$), the structural source of the precautionary motive.

5.2 Vignettes as belief shifts

Each information treatment exogenously shifts F while holding the state (a_0, b_0, y_0) fixed; the operational detail (groups, stages, randomization) is described in Section 2.2. The treatments fall into two blocks. The first varies *uncertainty*: the dispersion of a future-income distribution whose outcome probabilities are supplied to the respondent, so that the mean is held fixed by construction. The second introduces *ambiguity of uncertainty realization* by withdrawing those probabilities while keeping the support unchanged, so that the mean is no longer pinned down by design and the whole income distribution relies on the subjective distribution within respondents' minds.

Block 1: Income uncertainty (probabilities specified)

- *Benchmark* F_0: income is certain (guaranteed to remain unchanged), $y' = \bar{y}$.
- *uncertainty magnitude* $F_1(\delta)$: respondents are told that their household income over the next twelve months is equally likely to fall a fraction δ below or above its expected level $\bar{y}$, i.e. $y' \in \{\bar{y}(1-\delta),\, \bar{y}(1+\delta)\}$ with $\Pr = \{\frac{1}{2}, \frac{1}{2}\}$. This is a symmetric, mean-preserving

[7] Setting $R_a = R_b \equiv R$ admits interior solutions on both non-consumption margins; the non-trivial responses on both margins (Section 3) are consistent with a small wedge. The period-2 constraint is in Appendix 5.6.

spread: expected income is held fixed at $\mathbb{E}_{F_1}[y'] = \bar{y}$, so only the dispersion varies. The treatment assigns one of two dispersion levels, $\delta \in \{0.2, 0.4\}$—a ±20% or ±40% band around expected income—so the wider arm carries four times the variance of the narrower one, $\sigma^2(\delta) = (\delta\bar{y})^2$.

- *Path*: respondents reach a terminal δ either directly or via the other dispersion, so allocations at a given δ can be compared across routes.

Block 2: Ambiguity of uncertainty This block builds on the uncertainty design but withdraws the outcome probabilities, keeping the support $\{\bar{y}(1-\delta), \bar{y}(1+\delta)\}$ identical to the corresponding uncertainty arm. Because the probabilities are no longer supplied, the experimenter does not impose mean preservation: unlike $F_1(\delta)$, the mean of y' is not guaranteed to equal $\bar{y}$. Whether the household nonetheless behaves as if facing a mean-preserving spread depends on how it fills in the missing probabilities—an object of belief formation rather than of design. Section 5.4 formalizes the two candidate responses (ambiguity aversion versus retained probabilistic beliefs).

5.3 Precautionary premium and proportional response

Let z^0 and $z^*(\delta)$ be the optimal non-consumption allocations under F_0 and $F_1(\delta)$; the precautionary premium is $\psi(\delta) \equiv z^*(\delta) - z^0$. A second-order expansion of the Euler equation gives the Kimball approximation (Kimball, 1990),

$$\psi(\delta) \approx \tfrac{1}{2}\,\eta(c^0)\,\sigma^2(\delta) = \frac{\gamma+1}{2c^0}\,\delta^2\bar{y}^2, \tag{3}$$

with absolute prudence $\eta(c) = (\gamma+1)/c$ and $\sigma^2(\delta) = (\delta\bar{y})^2$. Allocations are recorded as budget shares, so the natural outcome is the proportional change in each share. Since $\Delta \log C \approx -\psi(\delta)/c^0$,

$$\boxed{\Delta \log C(\delta) \approx -\frac{\gamma+1}{2}\Big(\frac{\bar{y}}{c^0}\Big)^2\delta^2,} \tag{4}$$

with $\Delta \log S$ and $\Delta \log D$ positive multiples of $\psi(\delta)$ scaled by their baselines. The response scales with δ^2—variance, not standard deviation—and the proportional form is comparable across households with heterogeneous baselines.

5.4 Testable predictions

Equation (4) yields four predictions (Table 6), where $\pi^C_{20} \equiv \frac{1}{2}(\gamma+1)(\bar{y}/c^0)^2(0.2)^2$ is the percentage change of consumption response at $\delta = 0.2$ (analogously for S and D).[8]

*Block 1: Predictions under income uncertainty Predictions 1–3 characterize the precautionary response under known probabilities, where the premium $\psi(\delta)$ is well defined.

[8]The table reports percentage change responses, whereas the vignette imposes the budget restriction $C + S + D = 100$. Adding up holds exactly in levels ($\Delta C + \Delta S + \Delta D = 0$) but not in logs, so the three columns are linked through the baseline shares, $c^0\Delta \log C + s^0\Delta \log S + d^0\Delta \log D = 0$, rather than independent.

Table 6: Model Predictions: Proportional Response to uncertainty.

Treatment contrast	$\Delta \log C$	$\Delta \log S$	$\Delta \log D$	Ratio to π_{20}
$\delta : 0 \to 0.2$	$-\pi_{20}^{C}$	$+\pi_{20}^{S}$	$+\pi_{20}^{D}$	1
$\delta : 0 \to 0.4$	$-4\pi_{20}^{C}$	$+4\pi_{20}^{S}$	$+4\pi_{20}^{D}$	4
$\delta : 0.2 \to 0.4$	$-3\pi_{20}^{C}$	$+3\pi_{20}^{S}$	$+3\pi_{20}^{D}$	3
$\delta : 0.4 \to 0.2$	$+3\pi_{20}^{C}$	$-3\pi_{20}^{S}$	$-3\pi_{20}^{D}$	-3

- **Prediction 1 (Sign)**. $\pi^C < 0$, and $\pi^S, \pi^D > 0$.
- **Prediction 2 (Variance scaling)**. $\Delta \log X(\delta = 0.4) / \Delta \log X(\delta = 0.2) = 4$ for $X \in \{C, S, D\}$, invariant to whether the response is measured in levels, shares, or logs.
- **Prediction 3 (Path independence)**. Because $\psi(\delta)$ depends only on the terminal uncertainty level, cumulative responses from F_0 are state functions, invariant to the route by which δ is reached (Appendix 5.6). The test defines the path-dependence wedge

$$\begin{aligned} \Lambda_{20}^{X} &\equiv \Delta \log X\big[0 \to 0.4 \to 0.2\big] - \Delta \log X\big[0 \to 0.2\big], \\ \Lambda_{40}^{X} &\equiv \Delta \log X\big[0 \to 0.2 \to 0.4\big] - \Delta \log X\big[0 \to 0.4\big], \end{aligned} \tag{5}$$

  with the benchmark prediction $\Lambda_{20}^{X} = \Lambda_{40}^{X} = 0$. A nonzero wedge signals departures from the CRRA benchmark: $\Lambda_{20}^{C} < 0$ together with $\Lambda_{40}^{C} > 0$ indicates incomplete belief revision (anchoring to the stage-1 level), while an asymmetry $|\Lambda_{40}^{C}| \neq |\Lambda_{20}^{C}|$ that preserves the loss/gain coding of the two paths points to reference dependence. Since the vignettes hold (a_0, b_0, y_0) fixed, any wedge reflects beliefs or preferences, not mechanical state carryover.

*Block 2: Prediction under income uncertainty

- **Prediction 4 (uncertainty versus Knightian uncertainty)**. The two formalizations of the uncertainty treatment give opposite predictions. Under ambiguity aversion—maxmin expected utility (Gilboa and Schmeidler, 1989) or smooth ambiguity (Klibanoff et al., 2005)—the premium gains a component $\varphi(\delta) > 0$, so $\psi_{\text{Knight}} = \psi_{\text{uncertainty}} + \varphi$ and $\Delta \log X_{\text{Knight}} > \Delta \log X_{\text{uncertainty}}$, as in the household block of Ilut and Valchev (2023); Ilut et al. (2025). If instead households carry the 50–50 Stage-1 prior forward, F is unchanged and $\Delta \log X_{\text{Knight}} = \Delta \log X_{\text{uncertainty}}$.

5.5 Empirical counterparts

The choice variables map directly to the survey responses (Section 2.2): c to reported monthly expenditure, and the non-consumption allocation to self-reported saving s and debt repayment d, which lets the analysis separate precautionary saving from deleveraging. Subjective beliefs F enter through two conceptually distinct channels. *uncertainty*—dispersion under known probabilities—has both an observational and an experimental counterpart: it is measured from the probabilistic income questions (Section 2) via the income-normalized standard deviation $\tilde{\sigma} = \sigma / y$, and assigned exogenously by the magnitude and path vignettes (Section 2.2). *Knightian uncertainty*—dispersion when probabilities are withdrawn—has no $\tilde{\sigma}$ counterpart, since a

standard deviation presupposes a probability measure; it is identified only through the uncertainty vignette, which removes the probabilities while holding the support fixed.

5.6 Derivation of the Precautionary Premium

Budget and Euler equation. In period 2 the household consumes

$$c' = y' + R\,a' - R(R\,b_0 - x). \tag{6}$$

Combining (1)–(6), the optimal non-consumption allocation $z = a' + x$ satisfies the Euler equation

$$u'(c) = \beta R\,\mathbb{E}_F\big[u'(c')\big]\,. \tag{7}$$

Under the certainty benchmark F_0 this reduces to $u'(c^0) = \beta R\,u'(\bar{c}')$, defining z^0.

Second-order expansion. Expanding $u'(c')$ around $\bar{c}' = \mathbb{E}_F[c']$ and taking expectations,

$$\mathbb{E}_F\big[u'(c')\big] \approx u'(\bar{c}') + \tfrac{1}{2}\,u'''(\bar{c}')\,\sigma^2(\delta), \tag{8}$$

since the mean-preserving spread leaves $\bar{c}'$ unchanged. Substituting into (7) and comparing with the certainty case, the extra non-consumption required to restore the first-order condition is

$$\psi(\delta) \approx \tfrac{1}{2}\,\eta(c^0)\,\sigma^2(\delta), \qquad \eta(c) \equiv -\frac{u'''(c)}{u''(c)} = \frac{\gamma+1}{c}, \tag{9}$$

which with $\sigma^2(\delta) = (\delta\bar{y})^2$ gives (3). Dividing by c^0 yields the proportional form (4).

Path independence. Since $\psi(\delta)$ depends only on the terminal δ, cumulative responses measured from F_0 are state functions:

$$\Delta \log X\big[0 \to 0.4 \to 0.2\big] = \Delta \log X\big[0 \to 0.2\big] = \mp\,\pi^X_{20}, \tag{10}$$

$$\Delta \log X\big[0 \to 0.2 \to 0.4\big] = \Delta \log X\big[0 \to 0.4\big] = \mp\,4\pi^X_{20}, \tag{11}$$

for $X \in \{C, S, D\}$, with the upper sign for C and the lower for S, D. Equivalently, the two stage-2 increments are equal and opposite,

$$\Delta \log X(0.2 \to 0.4) = -\,\Delta \log X(0.4 \to 0.2) = \mp\,3\pi^X_{20}, \tag{12}$$

which is the single-step form of the restriction tested through the wedge (5).

6 Appendix: Data and Measurement

6.1 Detailed survey question

Income expectations probabilistic question In your opinion, how likely is it that the average monthly net income of your household will change as follows over the next twelve months?

Please allocate probabilities between 0 and 100 to each category, where 0 means that you are absolutely certain the event will not occur and 100 means that you are absolutely certain it will occur. The probabilities across all categories must sum to 100.

1. Decrease by €2,000 or more
2. Decrease by between €1,500 and less than €2,000
3. Decrease by between €1,000 and less than €1,500
4. Decrease by between €500 and less than €1,000
5. Decrease by between €250 and less than €500
6. Decrease by between €0 and less than €250
7. Increase by between €0 and less than €250
8. Increase by between €250 and less than €500
9. Increase by between €500 and less than €1,000
10. Increase by between €1,000 and less than €1,500
11. Increase by between €1,500 and less than €2,000
12. Increase by €2,000 or more

Note: The entered probabilities must add up to 100.

Expenditures in the previous Month You will now be shown some things that people can or need to spend money on in their everyday lives. If you think back to last month, roughly how much did you spend (in euro) on the following items?

Please provide an amount for each category. If you are not quite sure, please give a rough estimate.

1. Major purchases (e.g., car, furniture, electrical appliances)
2. Essential goods (e.g., food and beverages, non-food items such as cleaning products)
3. Clothing and footwear
4. Entertainment and recreation (e.g., restaurant visits, cultural events, gym)
5. Mobility (e.g., fuel, car loans and running costs, bus and train tickets)
6. Services (e.g., hairdresser, childcare, medical costs)
7. Travel and holidays
8. Housing costs (e.g., rent, mortgage, ancillary costs)
9. Savings (e.g., savings accounts, shares, bonds)

Planned expenditure over the next twelve months Are you likely to spend more or less over the next twelve months than in the last twelve months?

1. Plan to spend more
2. Plan to spend roughly the same

3. Plan to spend less

Note: The expenditure categories are the same as those listed in the previous question on last month's spending. Please select one response for each category.

Household monthly net income What is the total monthly net income of your household?

This refers to the total amount, including wages, salaries, income from self-employment, and pensions, in each case after deducting taxes and social security contributions. Please also include income received through public aid, earnings from rents and leases, housing allowances, child benefits, and any other sources of income. If your household income has varied over the last twelve months, please report the average of the last twelve months.

1. Less than €500
2. €500 to €999
3. €1,000 to €1,499
4. €1,500 to €1,999
5. €2,000 to €2,499
6. €2,500 to €2,999
7. €3,000 to €3,499
8. €3,500 to €3,999
9. €4,000 to €4,999
10. €5,000 to €5,999
11. €6,000 to €7,999
12. €8,000 to €9,999
13. €10,000 or more

Household balance sheet Components Respondents report the approximate value of the following asset and liability categories using ordered value brackets:

- Bank deposits
- Real estate (self-assessed market value)
- Securities (shares, bonds, including funds and ETFs)
- Equity in privately held businesses
- Other financial or real assets
- Outstanding mortgage debt
- Other outstanding loans (e.g., consumer credit, business loans, informal loans)

For each category, respondents select one of several increasing value intervals (including zero holdings). Detailed survey document in: https://www.bundesbank.de/en/bundesbank/research/survey-on-consumer-expectations-bop-hh/questionnaires-850746

6.2 Respondents' subjective income distribution

Probabilistic income expectations question. Perceived income uncertainty is measured using probabilistic survey questions that elicit respondents' subjective assessments of future household income changes over the next 12 months. Individuals report their perceived likeli-

hood of household income changes across income-change intervals,[9] allowing the construction of individual-level measures of expected income growth and expected income variability. The exact survey questions are reported in Appendix 6.1.

Eliciting Income uncertainty from Subjective Expectations. I measure perceived income uncertainty using probabilistic expectations questions. This choice differs from two common alternatives: realized income volatility and qualitative forecasts. Realized volatility reflects what happened, not what households believed in advance. Qualitative forecasts reveal direction but not magnitude. Probabilistic questions instead reveal each household's beliefs directly, without assuming these beliefs match the income process estimated by the econometrician (Manski, 2004, 2018).

Based on the probabilistic income expectations question, I construct a household-level measure of perceived income uncertainty as the standard deviation of the subjective distribution of expected income changes, following (Guiso et al., 1992). This measure captures the dispersion of each household's expectations: larger values indicate greater uncertainty about future income. Perceived uncertainty is direction-neutral. It may reflect downside uncertainty from potential income losses. It may also reflect upside uncertainty from possible income gains.

$$\sigma_{i,t} = \sqrt{\sum_{B=1}^{\text{Bin}} p_{i,t}^{(B)} \left(y_{i,t+1}^{(B)} - \mathbb{E}^*[y_{i,t+1}] \right)^2}. \tag{13}$$

Here, $p_{i,t}^{(B)}$ denotes the subjective probability assigned by household i at time t to income-change bin B, and $y_{i,t+1}^{(B)}$ denotes the midpoint of that bin. The subjective expected income change is given by

$$\mathbb{E}^*[y_{i,t+1}] = \sum_{B=1}^{\text{Bin}} p_{i,t}^{(B)} y_{i,t+1}^{(B)}. \tag{14}$$

To ensure comparability across households, I normalize perceived income uncertainty by current household income:

$$U_{i,t} \equiv \frac{\sigma_{i,t}}{Y_{i,t}}. \tag{15}$$

Similarly, perceived income growth is measured as expected income change relative to current household income:

$$G_{i,t} \equiv \frac{\mathbb{E}^*[y_{i,t+1}]}{Y_{i,t}}. \tag{16}$$

Unless otherwise stated, perceived income uncertainty refers to the income-normalized standard deviation $U_{i,t}$, while perceived income growth refers to the income-normalized subjective mean $G_{i,t}$.

[9]The intervals for the probabilistic question are: $>$ −2,000 euros, (−2,000, −1,500], (−1,500, −1,000], (−1,000, −500], (−500, −250], (−250, 0], (0, 250], (250, 500], (500, 1,000], (1,000, 1,500], (1,500, 2,000], and $>$ 2,000 euros. Negative values indicate income decreases and positive values indicate income increases.

Heterogeneity in subjective income distributions Figure 10 presents three representative subjective income distributions reported by respondents. The examples illustrate substantial heterogeneity in how households perceive future income uncertainty, in contrast to the common practice of imposing a single, homogeneous income process on all households. Some respondents place all probability mass on a single outcome, implying near certainty about future income. Most respondents spread probability mass across several positive-income states, while still others assign probability to both income gains and losses. These patterns show that subjective expectations are inherently distributional objects, not point forecasts, and that reducing them to a single assumed process discards meaningful variation across households.

I classify each individual using all observed panel responses into one of three categories. *Point mass* denotes individuals who report a single bin throughout the panel, implying near-certainty about future income. *Two-bin distribution* denotes individuals who report exactly two bins throughout the panel. *Beta* denotes individuals whose multi-bin responses are approximated by a beta distribution, the standard approach for smoothing discrete bin probabilities into a continuous subjective density.

Figure 11 shows the weighted share of individuals in each category. The large majority of respondents, 87.6 percent, report distributions well approximated by a beta distribution. A further 10.0 percent report a two-bin distribution, and only 2.4 percent report a point mass in every wave they are observed. This distribution of response types indicates that most households report genuine distributional beliefs rather than degenerate point forecasts, supporting the use of the full subjective distribution – rather than a single expected value – to construct perceived income uncertainty.

6.2.1 Summarizing Subjective Distributions by Moments

Having established that most individual subjective distributions are well approximated by a beta distribution, I next assess how much information is lost by summarizing each distribution using only its first two moments, expected income change ($\mu_{i,t}$) and income uncertainty ($\sigma_{i,t}$), rather than its full shape. Higher-moment shape differences (e.g., skewness, kurtosis) are limited for most respondents, so the economically relevant variation across individual beliefs is concentrated in these two moments.

Figure 12 classifies individuals by the sign of their expected income change (positive or negative) crossed with whether their perceived income uncertainty falls above or below the sample median (high or low uncertainty), alongside the point-mass and two-bin categories from Figure 11. Positive expectations paired with low uncertainty account for the largest share of individuals, at 33.1 percent, followed by positive expectations paired with high uncertainty (23.4 percent) and negative expectations paired with high uncertainty (20.4 percent). Negative expectations paired with low uncertainty account for the smallest share among the four beta-distribution categories, at 10.6 percent. This decomposition illustrates that the first and second moments jointly, rather than either moment alone, characterize the meaningful heterogeneity in subjective income beliefs across households.

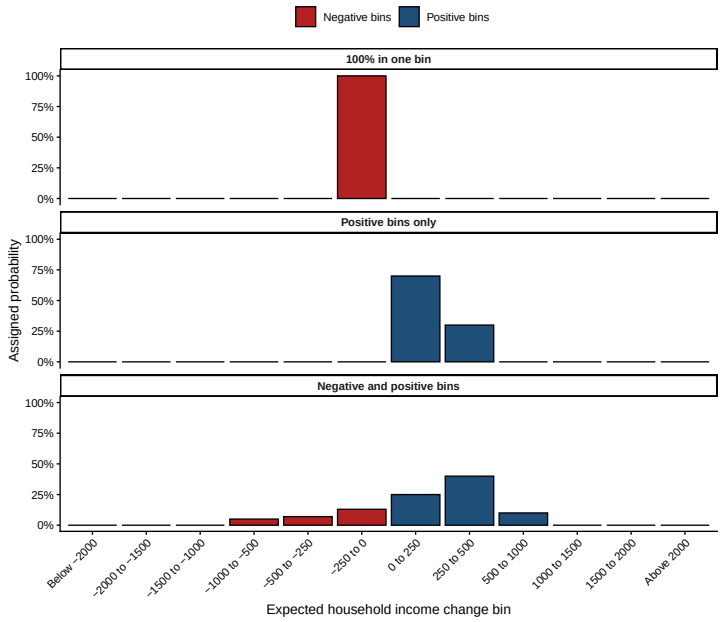


Figure 10: Heterogeneity in Subjective Income Distributions

Notes: The figure presents three representative subjective income distributions constructed from the probabilistic income expectations question. Respondents allocate probabilities across bins of future household income changes over the next 12 months. The examples illustrate the substantial heterogeneity in subjective income beliefs and show that households differ not only in the level of perceived income uncertainty but also in the shape of their subjective income distributions.

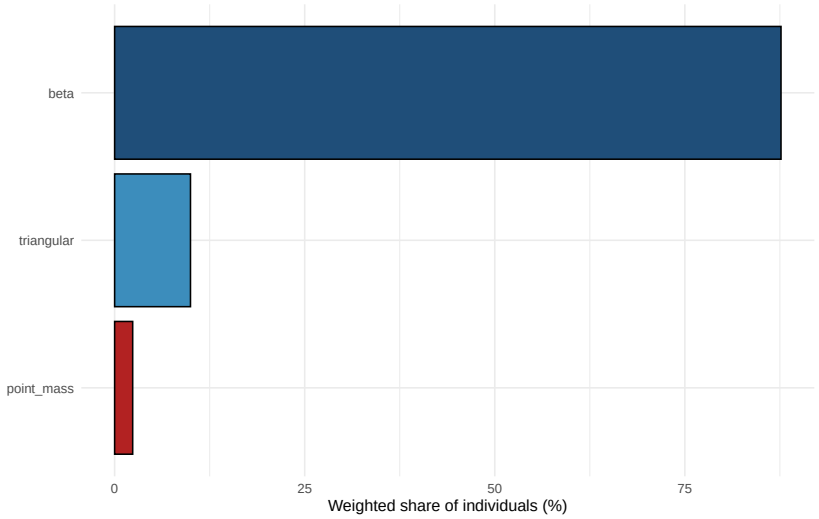


Figure 11: Distributional type at the individual level. Weighted shares: point mass 2.4%, two-bin distribution 10.0%, beta 87.6%.

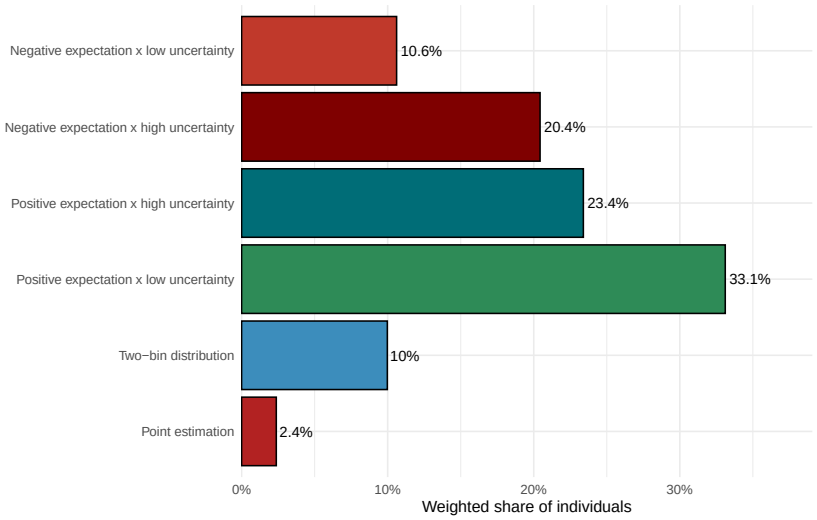


Figure 12: Subjective income distributions summarized by moments. Weighted shares of individuals by expectation sign and uncertainty level, together with point-mass and two-bin categories.

Aggregate distribution of perceived income uncertainty Pooling across households and time, most respondents perceive relatively small income uncertainty. Table 7 shows that the median household perceives income uncertainty equivalent to 2.8 percent of current monthly household income. Roughly one-quarter of the population reports near-zero uncertainty, placing almost all probability mass on a single expected outcome. Figure 13 shows that the overall distribution is right-skewed, with a long tail of households reporting much higher perceived uncertainty. This skewness is closely tied to employment status and occupation. Full-time employees and civil servants are concentrated near certainty, while self-employed households account for most of the mass in the right tail of high perceived uncertainty.

Co-moving between Perceived uncertainty and Perceived Growth Figure 13 also plots average perceived income growth conditional on each level of perceived income uncertainty, separately for households with positive and negative income outlooks. At any given level of perceived uncertainty, households split unevenly in the direction of their expectations: roughly two-thirds expect income growth, while the remaining one-third expect income declines. This pattern reveals an empirical challenge in observational expectations data. Perceived income uncertainty, the second moment of the subjective distribution, is systematically intertwined with perceived income growth, the first moment. The same measured level of uncertainty can therefore correspond to very different underlying beliefs, depending on whether it accompanies an optimistic or pessimistic outlook. This coupling makes it difficult to attribute observed consumption and saving responses to income uncertainty alone. It motivates an empirical strategy that can isolate the second moment and causally test its effect on financial decisions independent of the first.

Validating the subjective distribution I validate this measurement approach at the individual level, using each respondent's full panel history rather than pooled observations. Most respondents report genuine distributional beliefs: 87.6 percent of individuals report distributions well approximated by a beta distribution, and only 2.4 percent report a point mass throughout

Table 7: Distribution of Perceived Income uncertainty

	P5	P25	Median	P75	P95	Mean	SD	Obs.
Working age respondents	0.00	0.00	0.028	0.058	0.170	0.050	0.115	44,893

Notes: Perceived income uncertainty is measured as the income-normalized standard deviation of respondents' subjective income distributions. The distribution is calculated using the working age sample and is used to calibrate the experimental uncertainty shocks.

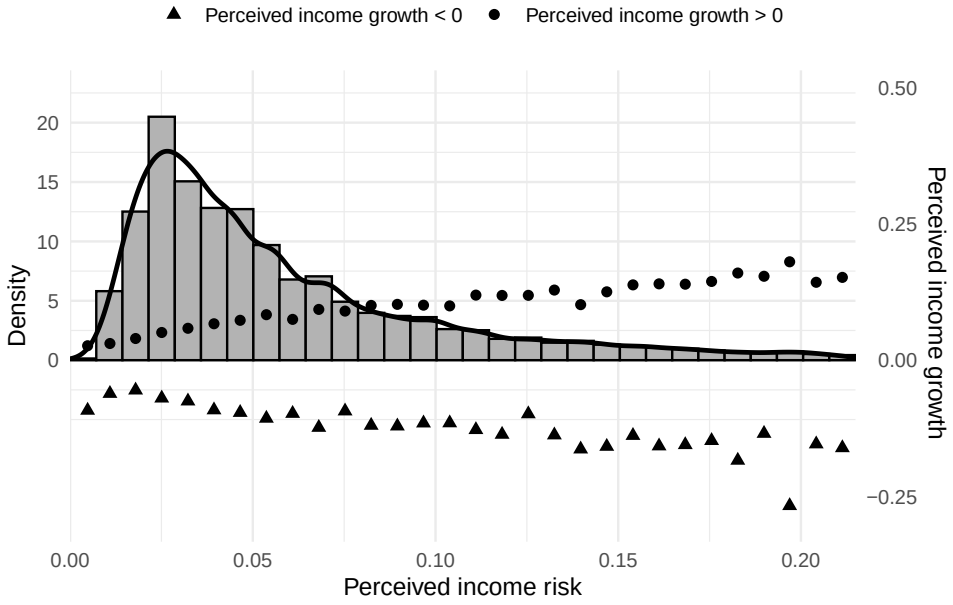


Figure 13: Perceived Income Growth Conditional on Perceived Income uncertainty

Notes: The histogram and solid line depict the distribution of perceived income uncertainty relative to household income, $U_{i,t} = \sigma_{i,t}/Y_{i,t}$. The markers plot average perceived income growth, $G_{i,t}$, conditional on each level of perceived income uncertainty. Circles indicate positive expected income growth, and triangles indicate negative expected income growth.

the panel (Appendix 6.2). These distributions are well summarized by two moments, expected income change and income uncertainty, since higher-moment shape differences are limited for most respondents (Appendix 6.2.1). Both moments vary substantially even within a single demographic cell: Figure 14 shows that full-time workers aged 35 to 44 range from near-zero perceived uncertainty to uncertainty exceeding 15 percent of household income, with expected income change spanning negative to strongly positive, and this dispersion is unrelated to gender. Demographic-cell proxies would therefore understate the true variation in perceived uncertainty and mute any heterogeneous response to income-uncertainty treatment operating through this channel.

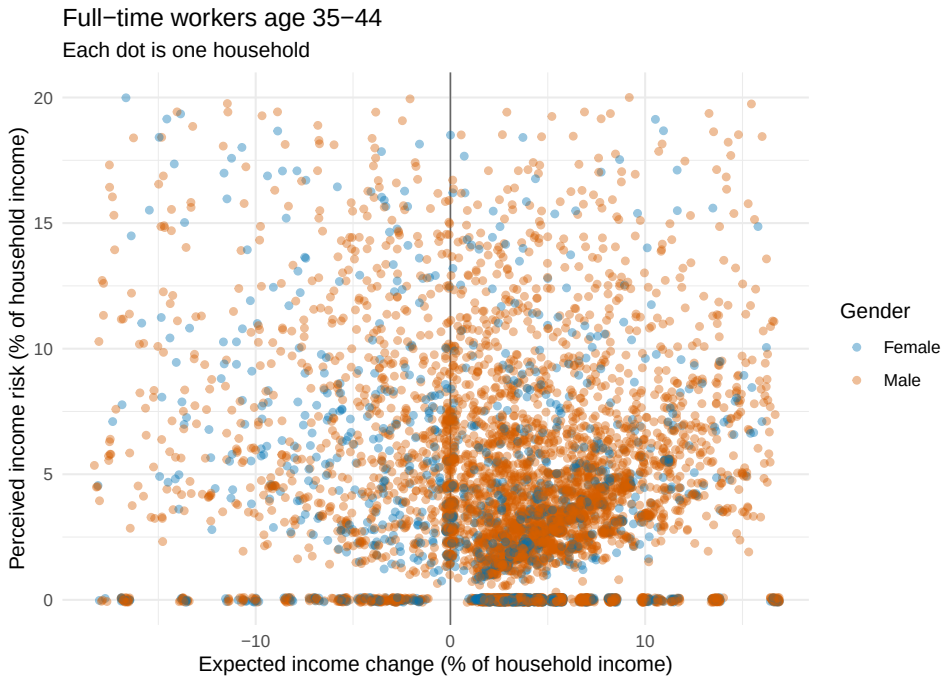


Figure 14: Perceived income uncertainty and expected income change, full-time workers aged 35–44.

Each dot is one household; color denotes gender. Perceived income uncertainty and expected income change are both expressed as a percentage of household income.

This appendix reports individual-level evidence supporting the construction of perceived income uncertainty described in Section 2. All statistics use each respondent's full sequence of panel responses, classified once per individual rather than pooled across household-wave observations.

6.2.2 Statistics summary

The sample is the working sample for wave November 2025 and is restricted to working-age individuals aged 20–65.

Table 8: Relationship Between Perceived Income uncertainty and Mean

uncertainty_bin	bin_mid	N_bin	N_pos	N_neg	exp_mean_pos	exp_mean_neg
[0,0.008]	0.005	66	61	5	0.028	-0.069
(0.008,0.016]	0.012	1367	1154	140	0.032	-0.058
(0.016,0.024]	0.020	3825	2971	493	0.042	-0.058
(0.024,0.032]	0.028	3745	2739	621	0.054	-0.071
(0.032,0.040]	0.036	3557	2381	637	0.063	-0.088
(0.040,0.048]	0.044	2773	1849	694	0.072	-0.097
(0.048,0.056]	0.052	2339	1438	628	0.079	-0.100
(0.056,0.064]	0.060	1591	1079	497	0.080	-0.103
(0.064,0.073]	0.068	1587	941	476	0.094	-0.120
(0.073,0.081]	0.077	1078	665	393	0.089	-0.097
(0.081,0.089]	0.085	902	549	343	0.105	-0.120
(0.089,0.097]	0.093	779	483	296	0.108	-0.128
(0.097,0.105]	0.101	883	402	317	0.107	-0.125
(0.105,0.113]	0.109	580	316	262	0.113	-0.122
(0.113,0.121]	0.117	473	268	202	0.123	-0.152
(0.121,0.129]	0.125	403	225	170	0.122	-0.101
(0.129,0.137]	0.133	368	205	156	0.133	-0.134
(0.137,0.145]	0.141	363	201	161	0.120	-0.175
(0.145,0.153]	0.149	287	162	125	0.131	-0.172
(0.153,0.161]	0.157	281	155	122	0.167	-0.154
(0.161,0.169]	0.165	284	119	111	0.128	-0.174
(0.169,0.177]	0.173	192	98	94	0.149	-0.149
(0.177,0.185]	0.181	172	90	82	0.176	-0.185
(0.185,0.193]	0.189	140	67	73	0.140	-0.143
(0.193,0.201]	0.197	173	81	89	0.230	-0.272
(0.201,0.209]	0.205	120	55	65	0.147	-0.160
(0.209,0.217]	0.213	99	56	40	0.126	-0.162
(0.217,0.226]	0.222	104	58	44	0.162	-0.155
(0.226,0.234]	0.230	82	43	35	0.205	-0.191
(0.234,0.242]	0.238	78	45	33	0.179	-0.199

Table 9: Summary Statistics: Working Sample

Variable	Mean	SD	P25	Median	P75	N
Demographics						
Age	43.01	13.08	32	43	55	2,665
Number of children	0.47	0.81	0	0	1	2,656
Male (%)	52.01					2,665
Partner (%)	61.20					349
Bachelor degree or above (%)	30.75					2,661
Employment						
Full-time (%)	61.16					2,665
Part-time (%)	16.95					2,665
Irregular (%)	1.78					2,665
Other employment (%)	18.80					2,665
Profession						
Salaried (%)	73.23					2,258
Civil servant (%)	9.47					2,258
Self-employed (%)	5.36					2,258
Other profession (%)	11.94					2,258
Income and Perceptions						
Monthly household income (€)	4,446	2,494	2,750	4,500	5,500	2,604
Perceived income growth	0.016	0.191	-0.010	0.030	0.061	841
Perceived income uncertainty	0.060	0.093	0.000	0.032	0.073	841
Household Balance Sheet						
Net wealth (€)	302,181	423,731	18,749	168,749	424,999	2,402
Cash on hand (€)	73,443	135,204	6,000	26,749	81,999	2,421
Liquidity-to-income ratio (%)	15.87	30.62	1.59	6.00	16.94	2,337
Debt-to-income ratio (%)	13.15	24.57	0.00	0.00	15.38	2,366
Budget Allocation						
Saving share	0.154	0.184	0.000	0.095	0.221	1,415
Consumption share	0.846	0.184	0.779	0.905	1.000	1,415

Notes: Statistics are computed using survey weights; N denotes unweighted observations. The sample is the working sample for wave November 2025 and is restricted to working-age individuals aged 20–65. *Perceived income growth* and *perceived income uncertainty* are measured relative to current monthly household income and can therefore be interpreted as the expected percentage change and percentage volatility of household income, respectively. *Liquidity-to-income* and *debt-to-income ratios* are calculated by dividing current liquid assets and outstanding debt by monthly household income. *Consumption share* is one minus the saving share. *Budget allocation* is measured at the monthly frequency, with major purchases reported separately. Partner status is based on family status reported in wave November 2025. Family status is not asked in every monthly wave.

7 Information treatment design

7.1 Current household income allocation

This appendix reports the full distribution of respondents' current household income allocation—elicited before any information treatment—and its determinants.

Table 10 reports summary statistics. On average, respondents allocate 65.84 percent of current household income to consumption, 20.79 percent to saving, and 13.38 percent to debt repayment; 1,342 of the 2,172 respondents report positive debt repayment in the prior question.

Table 10: Prior Allocation: Summary Statistics

Variable	N	Mean	P50	P75
Consumption	2172	65.84	70	80
Saving	2172	20.79	20	30
Debt repayment	2172[a]	13.38	5	25

Notes: Allocations are percentages and sum to 100.
[a] 1,342 respondents report positive debt repayment in the prior question; the remainder report zero.

Heterogeneity of current allocation Table 11 reports weighted regressions of prior-stage allocation shares on socio-demographic characteristics, belief measures, buffer-stock variables, and leverage measures. The results show that observable household characteristics are strongly related to allocation choices. Older respondents, partnered respondents, and households with children allocate a larger share to current consumption and a smaller share to saving. Higher household income is associated with a lower consumption share and a higher saving share, consistent with a stronger saving margin among higher-income households.

The belief channel plays a limited role once the full set of controls is included: neither perceived income growth nor perceived income uncertainty is statistically significant. In contrast, balance-sheet variables are strongly predictive. Higher cash-on-hand is associated with more saving and less debt repayment, while higher liquidity-to-income is associated with lower consumption and higher saving. Net wealth is significantly associated with all three allocation margins.

The leverage results suggest that debt composition is more informative than aggregate leverage. Total debt-to-income is not statistically significant conditional on the other controls, whereas secured debt and consumer debt are both associated with a substantially larger debt-repayment share and lower consumption and saving shares. The magnitudes are larger for secured debt, indicating that mortgage-related obligations are especially important for prior-stage repayment allocations.

7.2 Uncertainty and ambiguity preference

7.2.1 Preference collection question

Before the income-uncertainty treatments, respondents were randomly assigned to one of four groups and completed either an uncertainty-preference or an ambiguity-preference question. The

Table 11: Heterogeneity of current household income allocation

	Dependent variable		
	Consumption	Saving	Debt repayment
Socio-demographics			
Age	0.211***	-0.257***	0.046**
	(0.036)	(0.030)	(0.023)
Male	-0.502	0.730	-0.228
	(0.866)	(0.727)	(0.554)
Partner	2.554***	-3.891***	1.338**
	(0.977)	(0.820)	(0.625)
Number of children	1.388***	-1.826***	0.438
	(0.506)	(0.425)	(0.324)
Below bachelor	-2.272**	0.553	1.719***
	(0.909)	(0.763)	(0.581)
Full-time employment	-1.586	0.490	1.096*
	(0.985)	(0.827)	(0.630)
Salaried	-0.780	0.262	0.518
	(1.284)	(1.078)	(0.821)
Civil servant	-0.163	1.053	-0.889
	(1.807)	(1.517)	(1.156)
Self-employed	7.764***	-3.370*	-4.393***
	(2.113)	(1.774)	(1.351)
Log household income	-8.211***	9.436***	-1.225*
	(1.044)	(0.876)	(0.667)
Belief			
Perceived income growth	-0.006	0.015	-0.009
	(0.015)	(0.013)	(0.010)
Perceived income uncertainty	-0.024	0.064	-0.040
	(0.051)	(0.042)	(0.032)
Buffer-stock			
Log cash-on-hand	0.015	0.530***	-0.545***
	(0.204)	(0.171)	(0.130)
Liquidity-to-income	-0.095***	0.105***	-0.010
	(0.018)	(0.015)	(0.011)
Household net wealth	-0.665***	1.124***	-0.459***
	(0.175)	(0.147)	(0.112)
Leverage			
Debt-to-income	-0.005	0.030	-0.025
	(0.024)	(0.020)	(0.016)
Log secured debt	-12.728***	-8.331***	21.059***
	(1.267)	(1.064)	(0.810)
Log consumer debt	-3.642***	-2.978***	6.620***
	(1.203)	(1.010)	(0.770)
Observations	1,829	1,829	1,829
R^2	0.279	0.366	0.621
Adjusted R^2	0.268	0.356	0.615

Notes: The table reports weighted least-squares regressions. The dependent variables are prior-stage allocation shares for consumption, saving, and debt repayment. Right-hand-side variables are grouped into socio-demographics, belief, buffer-stock, and leverage channels. Household income enters in logs. Cash-on-hand, secured debt, and consumer debt enter as $\log(1+x)$, with nonpositive values set to zero before transformation. Net wealth enters using the inverse hyperbolic sine transformation, scaled by EUR 1,000, to accommodate zero and negative values. Belief variables are measured in percentage points. Missing right-hand-side variables are imputed using the respondent's latest nonmissing panel observation; imputation indicators are included but omitted from the table. Standard errors are in parentheses. Significance levels are denoted by *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

questions aim only to collect the respondents' preferences, not to elicit.

Uncertainty Preference (Groups 1 and 3) Respondents were asked to rank three possible income paths over the next twelve months:

> *"Please assume that you could choose how your monthly household income develops over the next twelve months. Three possible scenarios are shown below.*
>
> *Your monthly household income: 1. would be guaranteed to remain unchanged for the next 12 months. 2. would increase by 20% or decrease by 20%. Both scenarios would be equally likely. 3. would increase by 40% or decrease by 40%. Both scenarios would be equally likely. I do not have a preferred order."*

The order of the options was randomized across respondents. Participants were not required to provide a complete ranking and could indicate only their most preferred option if indifferent among the remaining alternatives. The task elicits ordinal preferences over mean-preserving spreads of different magnitudes under known probabilities.

Ambiguity preference (Groups 2 and 4) Respondents in Groups 2 and 4 faced a similar income-change support but evaluated probability information. Depending on randomization, income either increased or decreased by 20 percent (Group 2) or 40 percent (Group 4) over the next twelve months. The question read:

> *"Please assume that your monthly household income would either increase by 20% (40%) or decrease by 20% (40%) compared with your current household income. The change would only apply for the next 12 months.*
>
> *Which scenario would be preferable to you? 1. Both scenarios are equally likely (50–50). 2. The probability of the respective scenario occurring is unknown. 3. I have no preference."*

This design isolates attitudes toward ambiguity by varying whether probabilities are known or unknown while holding the income support fixed.

7.3 Balance across preference-response groups

After collecting preferences, I compare observable characteristics across risk- and ambiguity-preference response groups. All descriptive statistics are computed using survey weights. For a characteristic x_i, the reported group mean is

$$\bar{x}_g^w = \frac{\sum_{i \in g} w_i x_i}{\sum_{i \in g} w_i},$$

where w_i denotes the survey weight and g indexes the preference-response group.

To test whether characteristics differ across response groups, I use survey-weighted omnibus tests. For continuous variables, I estimate survey-weighted linear regressions of each characteristic on a full set of preference-group indicators and test the joint null that all group coefficients are equal to zero. This is the survey-weighted analogue of a between-group ANOVA test. For

categorical variables, I use Rao-Scott survey-weighted chi-square tests of independence between the characteristic and the preference-response group. I report both raw p-values and Benjamini-Hochberg adjusted p-values to account for multiple testing.

The results show limited evidence of systematic observable differences across preference-response groups. For risk preferences, perceived income risk is the strongest difference across groups: the raw omnibus test rejects equality across groups, although the Benjamini-Hochberg adjusted value is just above the 5 percent threshold. Other differences, including cash on hand, liquid assets-to-income, debt-to-income, and expected income change, are not statistically significant. For ambiguity preferences, none of the tested observable characteristics differ significantly across response groups after adjustment.

Table 12: Risk Preference Summary

Variable	Guarantee	20%	40%	No preference
Age	43.58	41.67	41.31	47.84
Female share	0.494	0.457	0.543	0.312
Full-time share	0.626	0.629	0.627	0.592
Part-time share	0.182	0.138	0.292	0.173
Casual share	0.015	0.053	0.034	0.000
Liquid assets-to-income ratio	16.27	17.77	19.37	12.12
Debt-to-income ratio	11.83	12.97	12.69	11.22
Cash on hand, euros	78,014	83,208	54,890	61,195
Perceived income risk	0.061	0.059	0.053	0.022
Expected income change	-0.003	-0.008	0.025	-0.032
Net wealth, euros	308,863	341,098	162,526	235,796
Observations	697	209	21	80

Notes: The table reports survey-weighted means by risk-preference response. Binary variables are reported as shares. Observations are unweighted. Weighted observations equal the sum of survey weights within each response group.

Table 13: Ambiguity Preference Summary

Variable	50–50	Unknown	No preference
Age	42.74	43.97	43.67
Female share	0.434	0.484	0.435
Full-time share	0.647	0.652	0.617
Part-time share	0.132	0.202	0.197
Casual share	0.006	0.000	0.000
Liquid assets-to-income ratio	17.06	14.87	19.47
Debt-to-income ratio	15.81	14.86	11.65
Cash on hand, euros	74,326	67,015	98,817
Perceived income risk	0.043	0.051	0.044
Expected income change	0.003	0.035	0.080
Net wealth, euros	360,058	272,115	336,325
Observations	236	366	296

Notes: The table reports survey-weighted means by ambiguity-preference response. Binary variables are reported as shares. Observations are unweighted. Weighted observations equal the sum of survey weights within each response group.

Table 14: Preference Group Tests

Variable	Risk p	Risk BH p	Ambiguity p	Ambiguity BH p
Age	0.077	0.172	0.745	0.840
Gender	0.086	0.172	0.557	0.836
Employment status	n.e.	n.e.	0.746	0.840
Liquid assets-to-income ratio	0.586	0.781	0.349	0.698
Debt-to-income ratio	0.947	0.947	0.267	0.698
Cash on hand	0.567	0.781	0.152	0.682
Perceived income risk	0.006	0.051	0.868	0.868
Expected income change	0.898	0.947	0.388	0.698
Net wealth	0.060	0.172	0.110	0.682

Notes: The table reports omnibus tests of equality across preference-response groups. For continuous variables, p-values come from survey-weighted linear regressions with preference group indicators and an omnibus F-test of joint equality across groups. For categorical variables, p-values come from Rao-Scott survey-weighted chi-square tests. BH p-values are adjusted using the Benjamini-Hochberg false discovery rate procedure. "n.e." denotes not estimable.

7.4 Baseline results

This appendix examines respondents' allocation response to the guaranteed-unchanged-income benchmark relative to their prior allocation, and whether that response reflects the benchmark being perceived as a negative income shock.

Table 15 reports paired within-person tests on log changes in allocation, $\Delta \log X_{i,G-prior} = \log(X_i^G) - \log(X_i^{prior})$. On average, consumption falls and saving rises moving from the prior to the guaranteed benchmark, while debt repayment is essentially unchanged.

Table 15: Average Response to the Guaranteed-Income Benchmark

$\Delta \log X_{i,G-prior}$	Consumption	Saving	Debt repayment
Guaranteed unchanged	-0.025^{***}	0.065^{***}	-0.001
	(0.003)	(0.008)	(0.006)
Obs.	2121	1898	1074

Method: Paired tests on within-person log changes, $\Delta \log X_{i,G-prior} = \log(X_i^G) - \log(X_i^{prior})$.

Why Is No Income Uncertainty Perceived as a Negative Shock? Figure 15 shows respondents' imputed expectations of income growth over the next twelve months: 68 percent expected positive growth, 27 percent expected a decline, and 4 percent expected no change. Because guaranteed, unchanged income is a shortfall relative to what most respondents expected, this asymmetry is consistent with the average response to the benchmark being negative. However, this does not compromise the benchmark's role in the design: its purpose is to fix perceived uncertainty and ambiguity at zero for every respondent, not to match each respondent's individual expected growth rate, and all subsequent uncertainty and ambiguity scenarios are measured against this common, uncertainty-free reference point.

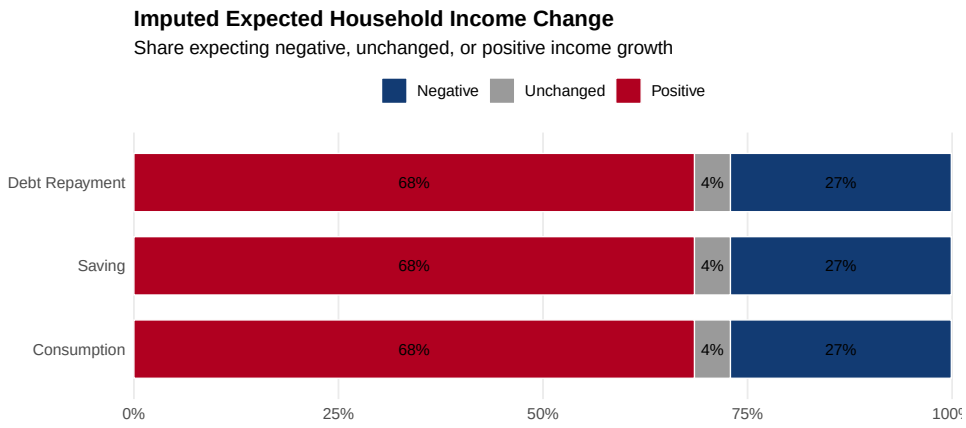


Figure 15: Imputed Expected Household Income Change

Notes: Bars show, separately for the consumption, saving, and debt repayment samples, the share of respondents whose imputed expected household income change over the next twelve months is negative, unchanged, or positive.

7.5 Statistical power calculation

The within-subject design concentrates statistical power in paired comparisons within arms rather than in comparisons across the full sample. Each treatment arm contains approximately 450–475 valid respondents, so comparisons that pool two arms sharing the same underlying comparison—such as the certainty-to-uncertainty comparisons, which combine Groups 1 with 2 and Groups 3 with 4—draw on roughly twice the sample and are correspondingly better powered than single-arm comparisons, such as the within-subject reversal or the uncertainty-to-ambiguity comparisons. Table 16 reports the standardized minimum detectable effect for each main comparison under a two-sided 5 percent test with 80 percent power. Pooled comparisons detect effects of approximately 0.09 standard deviations of the within-person change distribution, while single-arm comparisons detect effects of approximately 0.13 standard deviations. These magnitudes indicate the design is powered to detect economically modest reallocations across consumption, saving, and debt repayment.

Table 16: Minimum Detectable Effects for Main Within-Subject Comparisons

Comparison	Groups used	Approx. N	Standardized MDE
Certainty $\rightarrow$ $\pm$20% known uncertainty	1 and 2	900–950	0.09 SD
Certainty $\rightarrow$ $\pm$40% known uncertainty	3 and 4	900–950	0.09 SD
$\pm$20% $\rightarrow$ $\pm$40% known uncertainty	1	450–475	0.13 SD
$\pm$40% $\rightarrow$ $\pm$20% known uncertainty	3	450–475	0.13 SD
Known uncertainty $\rightarrow$ ambiguity at $\pm$20%	2	450–475	0.13 SD
Known uncertainty $\rightarrow$ ambiguity at $\pm$40%	4	450–475	0.13 SD

Notes: Standardized MDEs assume a two-sided 5 percent test and 80 percent power. For a paired within-subject comparison, the standardized minimum detectable effect is $(1.96+0.84)/\sqrt{N}$. To obtain outcome-specific minimum detectable effects in percentage points, multiply the reported standardized MDE by the empirical standard deviation of the within-person change in the relevant allocation share.

8 Robustness checks for stage-three Responses

Throughout this appendix, the robustness checks are conducted on the raw level of adjustment in each allocation rather than on the log changes used in the main text, so that the tests bear directly on the underlying responses.

8.1 Non-parametric Wilcoxon signed-rank test

The within-household responses for consumption, saving, and debt repayment contain many zero responses, are skewed, and include outliers, which can make parametric paired-t tests sensitive to distributional assumptions. The Wilcoxon signed-rank test instead asks whether the median response differs from zero under the null of a symmetric distribution around zero, without assuming normality.

Outcome	p-value		Zero-share		N
	20% shock	40% shock	20%	40%	
Consumption	<0.001***	<0.001***	0.534	0.391	961 / 938
Saving	0.002***	<0.001***	0.520	0.386	961 / 938
Debt repayment	0.001***	<0.001***	0.861	0.799	961 / 938

Table 17: Wilcoxon signed-rank test. H_0: symmetric distribution around zero. Zero-share = fraction with zero response.

Across all three outcomes and both shock sizes, the test rejects the null at conventional levels. Despite high zero-shares—particularly for debt repayment (0.861 at the 20 percent shock and 0.799 at the 40 percent shock)—the median responses remain significantly different from zero, confirming the parametric results reported in the main text.

8.2 Internal validity

A concern with the baseline estimates is that the treatment effects may reflect subgroup composition rather than the uncertainty treatment itself. To assess this, I estimate, for each outcome,

$$D_i = \alpha + \beta \cdot Group_i + \varepsilon_i,$$

where D_i is the household's treatment response and $Group_i$ is a demographic or financial characteristic, and test $H_0 : \beta = 0$, i.e. no systematic difference in responses across groups. Most subgroup coefficients are small and statistically insignificant, and the few significant estimates are limited and do not display a systematic pattern across outcomes. Treatment effects are thus broadly stable across demographic and financial characteristics, supporting internal validity.

8.3 Baseline dependence

Because responses are measured relative to each household's uncertainty allocation, the observed adjustments could in principle rescale baseline levels mechanically. To test this, I regress each uncertainty response on the corresponding baseline allocation and examine the share of variation explained. Although baseline allocations are statistically associated with the responses, they

explain less than 4 percent of the variation ($R^2 < 0.04$ across all three outcomes), which rules out purely mechanical rescaling from baseline levels.

8.4 Heterogeneity

Table 18 reports the regression underlying Figure 5: six separate weighted least-squares regressions, one for each combination of outcome (consumption, saving, debt repayment) and uncertainty treatment (20 percent, 40 percent). The dependent variable is the log difference between the uncertain and baseline allocation.

Socio-demographic effects are weak throughout. Age is significant in both treatments but economically small (0.001 to 0.005 log points per year). Civil-servant status raises saving by 24 log points in the 40 percent treatment ($p < 0.01$) with no counterpart elsewhere; self-employment lowers debt repayment by 8.5 log points in the 20 percent treatment ($p < 0.10$); log household income raises 40 percent debt repayment by 8.7 log points ($p < 0.10$). None of education, self-employment, or income is significant across more than one outcome or treatment, consistent with the balance-sheet margins carrying more of the heterogeneity than the demographic ones.

The buffer-stock panel shows the pattern is present but not monotone across measures. Middle-quartile cash-on-hand households cut consumption by 4.8 log points more and save 19.3 log points more than the bottom quartile in the 20 percent treatment (both $p < 0.05$), and top-quartile cash-on-hand households save a further 15.4 log points more ($p < 0.10$) — households with a cushion adjust their flows rather than draw down the cushion itself. Liquidity-to-income cuts the other way: middle-quartile households save 12.1 log points less than the bottom quartile ($p < 0.05$). Net wealth reverses sign again: middle-quartile net wealth households raise consumption by 3.3 log points and lower debt repayment by 5.2 log points (both $p < 0.05$), the opposite pattern from cash-on-hand. Top-quartile groups are rarely significant on any of the three measures, so the buffer-stock heterogeneity is concentrated in the middle of each distribution, not at the top.

Leverage gives the cleanest pattern in the 20 percent treatment. Middle-quartile debt-to-income households cut consumption by 8.8 log points and raise saving by 24.5 log points relative to the no-debt bottom quartile (both $p < 0.01$); top-quartile debt-to-income households cut consumption by 12.3 log points ($p < 0.01$), raise saving by 18.0 log points ($p < 0.10$), and raise debt repayment by 9.7 log points ($p < 0.10$). Holding secured debt complicates this: it raises consumption by 7.1 log points and lowers saving by 13.0 log points in the 20 percent treatment (both $p < 0.10$), the opposite sign from the debt-to-income results, then flips to lowering consumption by 6.2 log points and raising repayment by 12.5 log points in the 40 percent treatment (both $p < 0.05$). Consumer debt is null in the 20 percent treatment and only weakly negative for 40 percent consumption. The debt-to-income pattern itself disappears entirely at 40 percent.s

Table 18: Heterogeneity in Responses to Income Uncertainty: Main Specification

	20% uncertainty treatment			40% uncertainty treatment		
	Consumption	Saving	Debt repay.	Consumption	Saving	Debt repay.
Socio-demographics						
Age	0.001*	-0.001	0.000	0.003***	-0.005**	0.003*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Below bachelor	-0.017	0.012	-0.005	0.004	0.000	0.022
	(0.014)	(0.034)	(0.022)	(0.021)	(0.047)	(0.038)
Civil servant	0.002	0.001	0.054*	-0.038	0.241***	-0.019
	(0.021)	(0.052)	(0.033)	(0.030)	(0.069)	(0.051)
Self-employed	0.003	0.008	-0.085*	-0.011	0.033	-0.094
	(0.029)	(0.075)	(0.045)	(0.039)	(0.095)	(0.074)
Log household income	-0.017	0.034	0.007	0.013	-0.022	0.087*
	(0.016)	(0.043)	(0.027)	(0.019)	(0.048)	(0.046)
Buffer-stock						
Cash-on-hand: bottom 25%	0.008	-0.154*	0.039	0.040	0.063	0.137
	(0.036)	(0.091)	(0.058)	(0.049)	(0.113)	(0.086)
Cash-on-hand: middle 50%	-0.039	0.039	0.046	-0.024	0.068	0.163**
	(0.025)	(0.061)	(0.045)	(0.034)	(0.077)	(0.064)
Liquidity-to-income: middle 50%	0.019	-0.121**	0.014	0.036	-0.093	0.002
	(0.020)	(0.053)	(0.028)	(0.029)	(0.069)	(0.044)
Liquidity-to-income: top 25%	-0.012	-0.040	0.082	0.038	-0.121	0.171**
	(0.031)	(0.080)	(0.053)	(0.043)	(0.101)	(0.080)
Net wealth: middle 50%	0.033*	-0.070	-0.052*	0.033	-0.145**	-0.076
	(0.017)	(0.045)	(0.029)	(0.025)	(0.060)	(0.051)
Net wealth: top 25%	-0.034	-0.053	-0.002	0.000	-0.078	-0.015
	(0.023)	(0.058)	(0.040)	(0.033)	(0.077)	(0.061)
Leverage						
Debt-to-income: middle 50%	-0.088***	0.245***	0.035	0.022	-0.019	-0.061
	(0.028)	(0.075)	(0.051)	(0.042)	(0.099)	(0.096)
Debt-to-income: top 25%	-0.123***	0.180*	0.097*	0.013	0.055	-0.115
	(0.036)	(0.097)	(0.058)	(0.048)	(0.119)	(0.101)
Has secured debt	0.071**	-0.130*	-0.035	-0.064	-0.048	0.125**
	(0.029)	(0.077)	(0.036)	(0.039)	(0.096)	(0.056)
Has consumer debt	0.004	0.012	0.006	-0.062**	0.091	-0.031
	(0.023)	(0.060)	(0.028)	(0.032)	(0.075)	(0.043)
Constant	0.125	-0.159	-0.127	-0.310*	0.626	-0.918**
	(0.149)	(0.395)	(0.251)	(0.180)	(0.448)	(0.417)
Observations	815	718	433	805	678	423
R^2	0.080	0.064	0.077	0.044	0.094	0.080
Adjusted R^2	0.053	0.033	0.027	0.018	0.065	0.032

Notes: The table reports weighted least-squares regressions of Stage 1 responses to income uncertainty on household characteristics. The dependent variables are log differences between allocations under uncertain income and allocations under guaranteed income. Cash-on-hand, liquidity-to-income, net wealth, and debt-to-income are grouped into bottom quartile, middle two quartiles, and top quartile categories. For cash-on-hand, the top quartile is omitted; for liquidity-to-income, net wealth, and debt-to-income, the bottom quartile is omitted. Secured and consumer debt enter as indicators for positive debt. Missing right-hand-side variables are imputed using the respondent's latest nonmissing panel observation where available; imputation indicators are included but omitted from the table. Standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

9 Robustness checks for stage-four responses

As in Appendix 8, the checks below are conducted on the raw level of adjustment rather than on the log changes used in the main text. Let d_i^{inc} denote the within-subject change under the increase treatment ($20 \to 40$) and d_i^{dec} the change under the decrease treatment ($40 \to 20$).

9.1 Non-Parametric Wilcoxon Signed-Rank Test

The sequential responses contain many zero adjustments, skewness, and outliers, so I confirm the directional results without distributional assumptions. The test asks whether the median within-subject change differs from zero, H_0 : median$(d_i^{inc}) = 0$ (and analogously for d_i^{dec}); the final row additionally tests whether the absolute magnitudes are equal across directions, H_0 : median$(|d_i^{inc}|) =$ median$(|d_i^{dec}|)$.

Direction	Spending (a)	Saving (b)	Debt repayment (c)
Increase: 20%→40% ($n = 446$)	<0.001*** [zero=0.507]	0.004*** [zero=0.516]	0.002*** [zero=0.848]
Decrease: 40%→20% ($n = 473$)	<0.001*** [zero=0.531]	0.005*** [zero=0.550]	0.035** [zero=0.839]
Diff: $\lvert\uparrow\rvert - \lvert\downarrow\rvert$ (p-value)	0.401	0.456	0.665

Table 19: Wilcoxon signed-rank test (p-values). Zero-share = fraction with zero adjustment. Final row: test of symmetric magnitude (H_0: equal absolute adjustment).

Both directions are significant for all three outcomes ($p < 0.001$ to 0.035), and the absolute magnitudes of the upward and downward adjustments are statistically indistinguishable (all $p > 0.40$).

9.2 Order effect

A concern is that responses are driven by whether the large or small shock is presented first rather than by the uncertainty level itself. Comparing the escalating group ($G1 : 20\% \to 40\%$) and the de-escalating group ($G3 : 40\% \to 20\%$) at a common uncertainty level u, the null is $H_0 : E[Y_u \mid G1] = E[Y_u \mid G3]$, tested with a two-sample difference-in-means t-test.

Comparison	Consumption	Saving	Debt Repayment	N
40% shock: G1 stage 2 − G3 stage 1	0.709 (0.397)	−0.456 (0.574)	−0.253 (0.547)	446 vs 479
20% shock: G1 stage 1 − G3 stage 2	1.043* (0.053)	−0.675 (0.200)	−0.369 (0.195)	474 vs 473

Table 20: Mean difference (coefficient) with p-value in parentheses. Groups face the same uncertainty level in different orders. H_0: no order effect. $^*p < 0.10$.

Responses at a given uncertainty level do not depend on presentation order. The only marginal result is for spending in the 20 percent comparison ($p = 0.053$); saving and debt repayment show no order effect.

9.3 Sign test: direction of adjustment

Restricting attention to households with non-zero changes, the sign test asks whether positive and negative adjustments are equally likely, $H_0 : P(d_i > 0) = P(d_i < 0) = 0.5$, via a two-sided

binomial test $N_+ \sim \text{Binomial}(N_{\neq 0}, 0.5)$, conducted separately for each direction.

Direction		Consumption	Saving	Debt Repayment
Increase: 20%→40%	Share ↑	0.364*** ($n = 220$)	0.593*** ($n = 216$)	0.647** ($n = 68$)
Decrease: 40%→20%	Share ↓	0.671*** ($n = 222$)	0.371*** ($n = 213$)	0.408 ($n = 76$)

Table 21: Among participants who adjust, the share moving in the direction consistent with the shock. H_0: equal probability of increase vs. decrease. $^{**}p < 0.05$, $^{***}p < 0.01$.

Among adjusters the direction of change is overwhelmingly consistent with the shock—when uncertainty rises, spending falls (only 36.4 percent increase) while saving rises (59.3 percent), and the pattern reverses when uncertainty falls—confirming genuine behavioural responses rather than noise.

9.4 Magnitude (ratio) test

Finally, I check that the second-stage adjustment is economically meaningful relative to the initial allocation. Defining $R_i^{inc} = (Y_{40,i} - Y_{20,i})/|Y_{20,i}|$ and $R_i^{dec} = (Y_{20,i} - Y_{40,i})/|Y_{40,i}|$, the null is $H_0 : E[R_i^{inc}] = 0$ (and analogously for R_i^{dec}), tested with a one-sample t-test of the mean ratio against zero.

Direction		Spending (a)	Saving (b)	Debt (c)
Increase: 20%→40%	Ratio	−0.219** (−2.52)	0.178** (2.38)	0.292*** (3.11)
Decrease: 40%→20%	Ratio	0.156*** (3.91)	−0.131*** (−3.43)	−0.095 (−1.26)

Table 22: Ratio = adjustment relative to the allocation under the initially observed uncertainty level. t-statistic in parentheses. Signs are symmetric across directions. $^{**}p < 0.05$, $^{***}p < 0.01$.

The ratios are economically sizable and statistically significant in nearly every cell, and their signs reverse symmetrically between the increase and decrease directions, confirming that households make proportional, meaningful adjustments at each stage.

10 Full results for counter Effect under ambiguity

Tables 23–25 give the regressions behind Section 3.4: allocation on ambiguity preference, expected income change, perceived income uncertainty, demographics, and balance-sheet position, by uncertainty level and adjustment direction (any, decrease, increase, magnitude among adjusters).

Table 23: Predictors of Ambiguity Response: Consumption

	20% spread				40% spread			
Predictor	Pr(any)	Pr(dec.)	Pr(inc.)	Magn.	Pr(any)	Pr(dec.)	Pr(inc.)	Magn.
Prefer unknown prob.	-0.021 (0.054)	-0.023 (0.043)	0.002 (0.038)	0.118* (0.065)	-0.022 (0.064)	-0.053 (0.047)	0.031 (0.051)	-0.094* (0.048)
No preference	-0.004 (0.056)	0.004 (0.044)	-0.008 (0.039)	0.020 (0.063)	-0.046 (0.067)	-0.084* (0.049)	0.037 (0.053)	-0.109* (0.055)
Exp. income change	-0.023 (0.050)	-0.013 (0.040)	-0.010 (0.035)	0.034 (0.194)	-0.063 (0.145)	0.032 (0.105)	-0.094 (0.114)	-0.100 (0.128)
Perceived income uncertainty	0.730** (0.338)	0.547** (0.268)	0.183 (0.240)	0.032 (0.368)	0.009 (0.254)	-0.085 (0.184)	0.093 (0.200)	-0.002 (0.170)
Age	0.002 (0.002)	0.002 (0.002)	-0.001 (0.002)	-0.005** (0.002)	-0.000 (0.003)	-0.000 (0.002)	0.000 (0.002)	-0.001 (0.002)
Male	-0.034 (0.050)	-0.032 (0.040)	-0.003 (0.036)	0.035 (0.052)	-0.043 (0.058)	-0.021 (0.042)	-0.022 (0.045)	0.033 (0.042)
Below bachelor's	0.036 (0.047)	0.004 (0.037)	0.032 (0.033)	0.075 (0.051)	-0.095* (0.055)	-0.044 (0.040)	-0.051 (0.043)	0.007 (0.045)
Full-time employed	-0.016 (0.053)	0.002 (0.043)	-0.018 (0.038)	-0.014 (0.055)	-0.012 (0.060)	0.068 (0.044)	-0.079* (0.048)	0.068 (0.045)
Log hh income	0.056 (0.040)	0.021 (0.032)	0.035 (0.028)	0.002 (0.080)	-0.012 (0.065)	-0.003 (0.047)	-0.008 (0.051)	-0.011 (0.054)
Log cash-on-hand	0.009 (0.025)	0.011 (0.020)	-0.003 (0.018)	0.024 (0.036)	0.000 (0.032)	0.004 (0.023)	-0.003 (0.025)	0.065** (0.025)
IHS net wealth	-0.004 (0.007)	-0.003 (0.005)	-0.001 (0.005)	0.010 (0.007)	0.001 (0.007)	-0.004 (0.005)	0.005 (0.005)	-0.028*** (0.005)
Liquidity-to-income	-0.002 (0.001)	-0.001 (0.001)	-0.000 (0.001)	-0.002 (0.002)	0.001 (0.002)	-0.001 (0.001)	0.001 (0.001)	-0.001 (0.001)
Debt-to-income	-0.000 (0.001)	-0.000 (0.001)	0.000 (0.001)	0.001 (0.001)	-0.001 (0.001)	-0.000 (0.001)	-0.001 (0.001)	0.002** (0.001)
Observations	348	348	348	69	314	314	314	78

Notes: Coefficients are from linear probability models (direction outcomes: Pr(any adjustment), Pr(decrease), Pr(increase)) and OLS (Magnitude among adjusters) of the second-stage response to ambiguity, $\Delta_{\text{uncertainty},i} = \log(\text{unknown}_i) - \log(\text{known}_i)$, estimated separately by uncertainty level. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 24: Predictors of Ambiguity Response: Saving

	20% spread				40% spread			
Predictor	Pr(any)	Pr(dec.)	Pr(inc.)	Magn.	Pr(any)	Pr(dec.)	Pr(inc.)	Magn.
Prefer unknown prob.	-0.005 (0.059)	0.003 (0.041)	-0.008 (0.049)	0.079 (0.093)	-0.014 (0.070)	0.068 (0.055)	-0.081 (0.050)	0.231** (0.112)
No preference	0.006 (0.061)	-0.016 (0.042)	0.023 (0.050)	0.112 (0.095)	-0.049 (0.072)	0.016 (0.057)	-0.066 (0.051)	0.074 (0.132)
Exp. income change	-0.239 (0.178)	-0.056 (0.122)	-0.183 (0.146)	-0.443 (0.335)	-0.013 (0.161)	-0.042 (0.128)	0.029 (0.116)	-0.140 (0.295)
Perceived income uncertainty	0.263 (0.381)	-0.038 (0.262)	0.302 (0.313)	0.275 (0.613)	-0.206 (0.315)	-0.181 (0.250)	-0.025 (0.226)	0.411 (0.907)
Age	0.000 (0.002)	-0.002 (0.002)	0.002 (0.002)	-0.002 (0.004)	-0.003 (0.003)	-0.003 (0.002)	-0.000 (0.002)	0.003 (0.005)
Male	-0.041 (0.054)	-0.040 (0.037)	-0.001 (0.045)	0.061 (0.069)	-0.060 (0.062)	-0.013 (0.049)	-0.047 (0.044)	0.209* (0.107)
Below bachelor's	0.035 (0.050)	0.037 (0.035)	-0.002 (0.041)	0.048 (0.073)	-0.017 (0.060)	-0.032 (0.048)	0.015 (0.043)	0.078 (0.103)
Full-time employed	-0.042 (0.059)	-0.021 (0.040)	-0.021 (0.048)	0.055 (0.075)	-0.027 (0.064)	-0.088* (0.051)	0.061 (0.046)	0.204* (0.102)
Log hh income	0.071* (0.042)	0.039 (0.029)	0.032 (0.035)	0.200* (0.111)	0.038 (0.072)	0.007 (0.058)	0.031 (0.052)	-0.094 (0.124)
Log cash-on-hand	0.009 (0.027)	0.001 (0.019)	0.008 (0.022)	-0.126** (0.050)	-0.005 (0.036)	-0.005 (0.029)	-0.000 (0.026)	0.030 (0.055)
IHS net wealth	-0.004 (0.008)	-0.002 (0.005)	-0.002 (0.006)	0.011 (0.009)	-0.003 (0.008)	-0.003 (0.006)	0.001 (0.006)	0.013 (0.012)
Liquidity-to-income	-0.001 (0.001)	-0.000 (0.001)	-0.001 (0.001)	0.001 (0.003)	0.001 (0.002)	0.002 (0.001)	-0.000 (0.001)	-0.002 (0.002)
Debt-to-income	0.001 (0.001)	0.000 (0.001)	0.001 (0.001)	0.001 (0.001)	0.000 (0.001)	0.001 (0.001)	-0.001 (0.001)	-0.000 (0.002)
Observations	320	320	320	70	270	270	270	67

Notes: As before.

Table 25: Predictors of Ambiguity Response: Debt Repayment

	20% spread				40% spread			
Predictor	Pr(any)	Pr(dec.)	Pr(inc.)	Magn.	Pr(any)	Pr(dec.)	Pr(inc.)	Magn.
Prefer unknown prob.	-0.094 (0.071)	-0.044 (0.056)	-0.050 (0.047)	-0.292 (0.268)	-0.144* (0.078)	-0.113* (0.063)	-0.032 (0.054)	-0.002 (0.132)
No preference	-0.038 (0.072)	-0.010 (0.057)	-0.028 (0.047)	0.097 (0.204)	-0.110 (0.082)	-0.095 (0.066)	-0.015 (0.056)	-0.126 (0.159)
Exp. income change	-0.196 (0.358)	-0.396 (0.285)	0.200 (0.236)	1.476* (0.800)	-0.265 (0.225)	-0.059 (0.180)	-0.206 (0.154)	-0.859 (0.698)
Perceived income uncertainty	0.061 (0.544)	-0.017 (0.433)	0.077 (0.358)	-0.183 (2.588)	0.445 (0.437)	0.244 (0.349)	0.201 (0.299)	-0.692 (1.431)
Age	-0.001 (0.003)	-0.000 (0.002)	-0.000 (0.002)	-0.009 (0.011)	-0.001 (0.004)	0.004 (0.003)	-0.005* (0.003)	-0.001 (0.008)
Male	0.027 (0.073)	0.115** (0.058)	-0.088* (0.048)	-0.050 (0.229)	-0.074 (0.080)	-0.051 (0.064)	-0.023 (0.054)	0.200 (0.154)
Below bachelor's	0.049 (0.062)	0.017 (0.049)	0.033 (0.041)	0.044 (0.196)	0.001 (0.068)	-0.024 (0.054)	0.025 (0.046)	-0.252** (0.107)
Full-time employed	0.053 (0.079)	-0.004 (0.063)	0.058 (0.052)	0.061 (0.289)	0.050 (0.092)	0.021 (0.073)	0.029 (0.063)	-0.031 (0.147)
Log hh income	0.143 (0.087)	0.135* (0.070)	0.008 (0.057)	-0.293 (0.249)	0.007 (0.097)	0.044 (0.078)	-0.037 (0.067)	-0.251 (0.298)
Log cash-on-hand	-0.014 (0.035)	-0.031 (0.028)	0.017 (0.023)	0.113 (0.168)	0.060* (0.036)	0.026 (0.029)	0.034 (0.025)	0.178 (0.114)
IHS net wealth	-0.003 (0.007)	-0.001 (0.005)	-0.002 (0.004)	-0.029 (0.031)	-0.020*** (0.006)	-0.004 (0.005)	-0.016*** (0.004)	-0.006 (0.008)
Liquidity-to-income	0.003 (0.003)	0.003 (0.002)	-0.001 (0.002)	-0.002 (0.011)	-0.001 (0.002)	-0.000 (0.002)	-0.001 (0.002)	-0.006 (0.008)
Debt-to-income	0.000 (0.001)	0.001 (0.001)	-0.001 (0.001)	0.003 (0.004)	0.001 (0.001)	0.000 (0.001)	0.001 (0.001)	-0.003 (0.002)
Observations	174	174	174	26	153	153	153	27

Notes: As before.