Correlated Beliefs and Lifecycle Behavior *

Taha Choukhmane Samuel Earnest

March 2026

Abstract

Household subjective expectations are typically studied one domain at a time, but their effects on behavior depend on their joint distribution: for example, optimism about stock returns may raise equity shares, but if optimism extends to all assets, portfolio shares may remain unchanged. We measure this joint distribution using a new two-wave panel survey that elicits expectations about asset returns, labor market outcomes, inflation, and life expectancy for a representative sample of U.S. adults. We find that beliefs co-move along a small number of common factors, sort individuals into psychologically interpretable types, and display substantial within-person persistence once measurement error is corrected for. We incorporate this belief structure into a lifecycle model of consumption, saving, portfolio choice, and housing to quantify the behavioral and welfare implications of correlated expectations. Relative to a benchmark with equally dispersed but uncorrelated beliefs, the observed correlation structure substantially attenuates the pass-through from any single expectation to behavior and reduces the average welfare cost of belief heterogeneity by nearly one-third. Studying expectations one domain at a time thus overstates the behavioral and welfare consequences of belief heterogeneity.

*Choukhmane is with MIT Sloan School of Management; Earnest is with the MIT Sloan School of Management. We are thankful for comments from Christopher Palmer, Tong Liu, Jonathan Parker, and David Thesmar. We are also thankful for superb RA work from Matthew Akuzawa. Emails: tahac@mit.edu and searnest@mit.edu.

1 Introduction

A household that is too optimistic about its future earnings might save too little. But suppose the same household is also too optimistic about its longevity. The two effects push in opposite directions, and the household might end up saving the right amount. This possibility, that economic expectations interact and potentially offset each other, points to the limits of studying households' subjective expectations one at a time. There is now substantial evidence that heterogeneous beliefs about life expectancy (Heimer et al., 2019), income growth (Rozsypal and Schlafmann, 2023), employment prospects (Mueller et al., 2021; Caplin et al., 2023), stock returns (Giglio et al., 2021; Velásquez-Giraldo, 2024), home prices (Liu and Palmer, 2023), and inflation (D'Acunto et al., 2024) each matter for households' saving, portfolio choice, and home ownership decisions. How these beliefs interact with each other depends on their joint distribution, an object that has yet to be measured systematically.

In this paper, we study the joint distribution of household subjective expectations and its implications for lifecycle financial behavior. To do so, we fielded two rounds of an online survey to a representative U.S. sample, eliciting expectations across ten economic domains from the same individuals six months apart. Our survey covers the key outcomes that matter for lifecycle consumption, saving, and investing decisions: asset returns, home prices, income growth, employment prospects, inflation, interest rates, and survival probabilities. We then use this new data to describe the joint distribution of household expectations and to calibrate a rich lifecycle model with multiple assets, housing, and labor market dynamics. Using the model, we (i) quantify how belief correlations affect, and in most cases dampen, the pass-through from any single expectation to behavior, and (ii) examine the lifetime welfare implications of correlated beliefs.

We begin by documenting three empirical facts that characterize the survey evidence and motivate our modeling decisions. First, beliefs co-move systematically across individuals and can be summarized by just three factors that account for 53% of the variation across ten distinct domains. Using the factor analysis approach of Stango and Zinman (2023), we identify an *asset*

factor (loading on stock returns and house price growth), a *rates factor* (loading on interest rates and inflation), and a *human capital factor* (loading on wages, employment prospects, and life expectancy).

Second, individuals cluster into interpretable belief types that correlate with broader measures of psychological optimism. Using k-means clustering, we identify three distinct groups: *General Pessimists* (37%), who hold consistently negative views across all domains; *Asset Optimists* (21%), who are bullish on financial returns and housing but moderate about their own labor market prospects; and *Human Capital Optimists* (42%), who expect strong wages and employment outcomes but low asset returns. These types align with psychological measures of optimism (Scheier and Carver, 1985), which were excluded from both the factor analysis and clustering. That these independent measures correlate with our clusters suggests that individual psychology may be an important source of belief heterogeneity.

Third, beliefs are highly persistent within individuals, though raw cross-round correlations understate this persistence due to measurement error. Across two survey rounds six months apart, a majority of respondents remain in the same tercile of each of the three belief factors. Exploiting a re-elicitation design that asks the same expectations twice within each survey session, we estimate that 63% of observed cross-round variation reflects noise rather than genuine belief updating, implying a cross-round correlation of 0.58 after correcting for measurement error compared to a raw correlation of 0.34.

What do these empirical patterns imply for household decisions and welfare? We address these questions using a rich lifecycle model. Our focus is on the heterogeneity and correlations between beliefs, rather than the average level of beliefs. We therefore assume that the population holds correct beliefs on average, while allowing individual beliefs to be heterogeneous and calibrating their dispersion, correlations, and dynamics to match our survey responses. Agents make optimal consumption, saving, portfolio, and home ownership decisions using their subjective beliefs, but face realizations drawn from homogeneous processes of asset returns, labor prospects, and survival.

Using this model, we find that ignoring the correlations between beliefs overstates both the behavioral and welfare consequences of belief heterogeneity. On the behavioral side, taking belief correlations into account substantially attenuates the pass-through from any single expectation to behavior. We show this by comparing simulations in which beliefs are distorted one at a time, as is standard in the literature, against simulations calibrated to the full correlated belief structure. The equity share response to a one percentage point increase in expected stock returns falls by 21 percent, with even more pronounced attenuation for bond shares and consumption. On the welfare side, holding fixed the mean and dispersion of each belief, replacing uncorrelated with empirically correlated belief assignments reduces the average welfare cost of belief heterogeneity by nearly one-third, with an attenuation that is uniform across the entire distribution of welfare losses.

1.1 Related Literature

This paper contributes to several strands of literature. First, it relates to the extensive research on subjective household expectations, which documents substantial heterogeneity and systematic departures from full-information rational expectations. Prior work has examined expectations across a wide range of domains, including life expectancy (Hurd and McGarry, 1995), employment (Mueller et al., 2021; Caplin et al., 2023), stock returns (Giglio et al., 2021), inflation (D'Acunto et al., 2024), and home prices (Armona et al., 2019; Liu and Palmer, 2023). We contribute to this literature by analyzing the joint distribution of beliefs across domains rather than studying each in isolation.

Second, this paper relates to a growing literature on correlations among behavioral biases and on the transmission of beliefs to economic behavior. On the correlation side, Stango and Zinman (2023) and Chapman et al. (2023) study the correlation structure of preference-related biases; we extend this line of inquiry to subjective expectations and integrate the resulting correlation structure into a quantitative lifecycle model. On the transmission side, Giglio et al. (2021) find that the sensitivity of investment choices to beliefs is small on average, attribut-

ing these low elasticities to measurement error and trading frictions. Liu and Palmer (2026) show in the context of housing that survey-reported expectations may not fully capture the beliefs that enter decision-making, while Enke et al. (2024) show that cognitive uncertainty can further attenuate the pass-through from beliefs to behavior. We identify a novel attenuation channel: when beliefs are correlated across domains, the sensitivity of any single behavior to any single belief is systematically dampened.

Finally, we contribute to the growing literature that combines models of economic behavior with custom surveys (Bachmann et al., 2022; Fuster and Zafar, 2023; Stantcheva, 2023; Caplin, 2025). Survey responses bring rich variation to identify heterogeneity that is generally missing from administrative datasets but critical for model calibration (see Koşar and O'Dea, 2023 for a review). Ameriks et al. (2020) and Indarte et al. (2025) use surveys to calibrate key parameters in consumption models; and Carranza et al. (2026) use hypothetical survey choices to calibrate a model of retirement contributions. Closest to our paper, recent work has incorporated survey-measured beliefs into calibrated lifecycle models, typically focusing on a single expectation domain, such as stock return beliefs (Velásquez-Giraldo, 2024), income expectations (Rozsypal and Schlafmann, 2023), or life expectancy beliefs (Heimer et al., 2019; O'Dea and Sturrock, 2023). We build on this literature by incorporating the full joint distribution of beliefs across multiple domains, including their correlation structure, into a life-cycle model and showing that correlations are an important determinant of how belief heterogeneity maps into behavior and welfare.

The remainder of the paper is organized as follows. Section 2 presents a simple theoretical framework that motivates the measurement of joint beliefs. Section 3 describes the survey design. Section 4 documents empirical findings on the joint structure of beliefs across domains. Section 5 introduces the lifecycle model. Section 6 quantifies how belief correlations shape the impact of belief heterogeneity on lifecycle behavior and welfare. Section 7 concludes.

2 Theoretical Framework

To clarify which expectations govern lifecycle saving and portfolio decisions, we first present a simple theoretical framework. It guides both our survey design (i.e., determining which expectations to elicit) and the specification of the lifecycle model we later estimate. The key insight is that many combinations of expectations generate observationally equivalent behavior, motivating our focus on the joint distribution rather than any single expectation in isolation.

2.1 Setup

Consider a household i with CRRA preferences and elasticity of intertemporal substitution EIS $= 1/\sigma$ who maximizes lifetime utility subject to an intertemporal budget constraint. In each period t, the household receives income y and chooses consumption c_{it}. Any remaining wealth is saved and earns the real risk-free return $r - i$, where r is the nominal return and i is the inflation rate. The household works until period A_r and retires, living until death at period A. For clarity we take income to be flat over the working life and relax this in the full lifecycle model in Section 5. Formally, the household solves

$$\begin{aligned} \max_{\{c_{it}\}} \quad & \sum_{t=0}^{A-1} \beta^t \frac{c_{it}^{1-\sigma}}{1-\sigma} \\ \text{s.t.} \quad & \sum_{t=0}^{A-1} \frac{c_{it}}{(1+r-i)^t} = \sum_{t=0}^{A_r-1} \frac{y}{(1+r-i)^t}. \end{aligned} \tag{1}$$

2.2 Belief Heterogeneity

We allow household i to hold heterogeneous beliefs about the key parameters governing intertemporal choice: future income, life expectancy, the nominal return, and inflation. We parameterize these as multiplicative deviations from the true values,

$$\widetilde{y}_i = (1+\epsilon_{y,i})\, y, \qquad \widetilde{A}_i = (1+\epsilon_{A,i})A, \qquad \widetilde{r}_i = (1+\epsilon_{r,i})\, r, \qquad \widetilde{i}_i = (1+\epsilon_{i,i})\, i,$$

where a positive ϵ indicates optimism relative to the true parameter. The household chooses consumption optimally under these subjective beliefs, yielding initial consumption $\widetilde{c}_{i0}$. A first-order approximation of the percentage deviation from unbiased consumption c_0 gives

$$\frac{\widetilde{c}_{i0} - c_0}{c_0} \approx \underbrace{\epsilon_{y,i}}_{\text{income}} - \underbrace{\phi_A A \epsilon_{A,i}}_{\text{longevity}} - \underbrace{\left(\phi^r_{PI} + \phi_r(\text{EIS} - 1)\right)\left(r\,\epsilon_{r,i} - i\,\epsilon_{i,i}\right)}_{\text{real return}}, \tag{2}$$

where $\phi_A > 0$ captures sensitivity to life expectancy (increasing in A and A_r), $\phi^r_{PI} > 0$ captures the response of permanent income to real interest rates, and $\phi_r(\text{EIS} - 1)$ captures the net effect of intertemporal substitution and wealth effects on consumption. Note that when EIS > 1, the substitution effect dominates and higher perceived real returns reduce consumption. When EIS < 1, the wealth effect dominates and higher perceived real returns raise it. The combined term $r\,\epsilon_{r,i} - i\,\epsilon_{i,i}$ reflects that what matters for consumption choices is the perceived *real* return. In this sense expecting high nominal returns and high inflation have partially offsetting effects. Equation (2) delivers three implications that directly shape our empirical approach.

Multiple beliefs enter simultaneously. Income optimism ($\epsilon_{y,i} > 0$) raises consumption by expanding perceived lifetime wealth, whereas longevity optimism ($\epsilon_{A,i} > 0$) lowers it by lengthening the effective planning horizon. Misperceptions of the real return affect consumption through both the permanent income and intertemporal substitution channels. No single belief governs consumption in isolation.

Belief across domains can offset one another. An individual who is optimistic about both income and longevity experiences partially canceling effects on consumption as the income effect pushes up and the longevity effect pushes down. Similarly, optimism about nominal returns may be offset by optimism about inflation. As a result, households with very different belief profiles may exhibit similar observed behavior, and households with similar belief profiles may behave differently depending on whether their beliefs reinforce or offset one another. This

implies that the joint distribution of beliefs, not its marginal components is an economically relevant object.

Heterogeneous beliefs can lead to observationally equivalent behavior. A household that consumes aggressively may be optimistic about income, pessimistic about longevity, optimistic about real returns, or some combination of all three. This observational equivalence means that reduced-form regressions of behavior on a single belief will be misspecified if other beliefs are omitted, even if those beliefs are independently measured. Correcting for this requires simultaneously measuring expectations across all relevant domains and modeling their joint influence on behavior, which is precisely what our survey and lifecycle framework are designed to do.

2.3 Portfolio Choice

The same logic extends to portfolio choice. In the canonical mean-variance framework, the optimal risky asset share depends on perceived excess returns, perceived variances, and perceived covariances across asset classes. Because optimism about stocks may be partially offset by optimism about bonds or real estate, what drives portfolio allocations is the joint distribution of beliefs across assets, not any single expectation in isolation. A researcher who observes only stock return beliefs and regresses portfolio shares on them will therefore recover an attenuated coefficient. Unmeasured beliefs about other asset classes introduce a classic omitted variable bias, since these beliefs belong in the regression and are correlated with the included regressor.

These simple exercises motivate our survey and modeling approach. Consumption and portfolio choice depend on the joint distribution of subjective expectations across income, employment, longevity, asset returns, interest rates, and inflation, motivating a survey that elicits all of these for the same individuals rather than studying any one in isolation.

3 Survey Data

3.1 Survey Design

Studying belief correlations requires observing expectations across multiple domains for a representative sample, a combination that no existing public dataset provides. For example, the Survey of Consumer Expectations covers stock returns and inflation but not life expectancy or borrowing rates. The Health and Retirement Study includes life expectancy and equity beliefs but not risk-free rates and wage growth, and is restricted to individuals over 50.

To overcome this challenge, we ran a representative survey eliciting subjective expectations over ten different relevant outcomes. The survey was administered on Prolific, an online data collection platform that has been widely used and validated in prior academic research (Palan and Schitter, 2018; Bergman et al., 2020; Gorodnichenko and Yin, 2024). We choose Prolific for three reasons. First, Prolific verifies each user on its platform requiring them to provide a photo ID accompanied with a selfie video ensuring that each account is associated with an actual human being. This significantly decreases the likelihood of bots taking surveys on their platform. They additionally include AI-behavior checks in the onboarding process and continuous in-app monitoring that detect and penalize AI-assisted responses. Second, Prolific pays their participants for each completed survey. Compensation helps ensure that participants are willing to engage with the survey seriously, leads to higher completion rates, and can increase the quality of responses (Bergman et al., 2020). In our survey, we compensate respondents at a rate of approximately $12 per hour and additionally include both a comprehension check and an attention check. Lastly, Prolific offers the ability to survey individuals from a sample matched to the U.S. census age, gender, and racial composition.

Our survey was run in two rounds. The first round was conducted in July 2025 on approximately 1,000 U.S. adults. The second round was conducted in January 2026 (approximately 6-months later) and contains an additional 1,000 responses. We attempted to re-survey as many individuals from the July 2025 survey as possible in our second wave in January 2026.

We successfully collected responses from 600 of the original 1,000 individuals in the second survey round and fill the remaining 400 slots with individuals whose demographics match those who did not participate in the second wave of the survey.

3.2 Expectation Elicitation

The survey consists of three main sections. First, we collected demographic and financial information not already provided by Prolific. For example, we ask them for their total personal income before taxes during the past 12 months, the share of their retirement account/Roth IRA/401(k) that has been placed in equities, and several financial literacy questions. We also elicit psychological traits, including an index of dispositional optimism based on Scheier and Carver (1985)

Next, individuals are asked to elicit their expectations for ten key outcomes relevant for household saving and investment decisions: stock returns, local home price growth, national home price growth, Treasury bill rates, mortgage rates, inflation, wage growth, life expectancy, the probability of job loss, and the probability of finding a job conditional on job loss. Expectations were elicited as point forecasts using a continuous slider interface. To fix ideas, when eliciting stock return expectations, respondents were asked:

> *"Imagine that, as a long-term investment, you invest in the S&P 500 stock market index fund. If you held that fund for the next 20 years, what average annual return would you expect?"*

This framing is designed to capture beliefs relevant for life-cycle planning while abstracting from short-run fluctuations and market timing. We randomize the order of the belief elicitation questions.

Lastly, we ask individuals their beliefs about asset return correlations. For example, respondents are asked

"If the SP 500 stock market index were to fall by 10% over the next year, what do you think would happen to average U.S. home prices over the same year?"

Respondents are then offered a set of multiple choice questions ranging from *"home prices fall by more than 10%"* to *"home prices rise by more than 10%"*. We included these questions only in round 2. The full survey instrument can be found in Figure A1.

The biggest concern when eliciting expectations in a survey is measurement error. Point estimate elicitation can be particularly subject to measurement error. It has been shown that asking respondents to report expectations in the form of a point estimate can induce positive bias (Hartzmark and Sussman, 2024). If we were explicitly interested in the level of expectations individuals hold this would be problematic. However, because we are studying the correlation, not the level, of household expectations this is less of an issue. As long as the magnitude of positive bias for each expectation is roughly the same across individuals the correlations between biased expectations and unbiased expectations should be approximately equal. [1] In addition to using point estimate elicitations there are a number of other ways measurement error may impact responses. To explore these we ask respondents to re-elicit their stock return expectation at the end of each survey round. Section 4 uses these re-elicitations to explore the extent to which our survey contains measurement error.

3.3 Descriptive Statistics

Table 1 presents demographics and summary statistics for the ten primary expectations elicited in our survey for the full set of respondents (left) and the 600 individuals who responded to both survey rounds (right). Across all domains, ranging from macro-level variables like inflation to personal outcomes like job security, we find significant dispersion in beliefs. For example, while the median respondent perceives a 70% probability of living to age 75, the

[1] In pilot surveys we additionally tried eliciting distributions of beliefs as recommended by Hartzmark and Sussman (2024). We found these generally lined up with point forecasts but lead to longer survey duration, inattention, and respondent confusion. Given we are focused on correlations, not levels we chose to use point forecasts.

Table 1: Summary Statistics

	All Respondents			Repeat Respondents		
	Mean	S.D.	P50	Mean	S.D.	P50
Panel A: Expectations						
Stock returns	12.31	7.01	10.10	11.92	6.82	10.00
Home prices (national)	7.27	4.10	6.50	7.32	4.02	6.60
Home prices (local)	8.98	5.55	8.10	8.79	5.49	7.90
Inflation	4.47	2.54	3.70	4.30	2.40	3.50
Mortgage Rates	8.00	3.85	6.80	7.78	3.70	6.60
T-bill Returns	5.51	3.71	4.30	5.28	3.53	4.10
Wage growth	4.89	4.55	3.90	4.55	4.39	3.40
Job Loss Prob.	24.16	22.61	16.00	23.11	22.22	15.00
Job Finding Prob.	50.94	29.94	50.00	49.34	30.37	50.00
Prob. Live to 75	66.10	23.27	70.00	66.20	23.36	70.00
Panel B: Subjective Correlations						
Corr. Stock/Tbill	0.04	0.14	0.00	0.04	0.15	0.00
Corr. Stock/Inflation	-0.16	0.53	0.00	-0.15	0.51	0.00
Corr. Stock/Wages	0.13	0.36	0.00	0.11	0.35	0.00
Corr. Stock/Housing	0.09	0.57	0.00	0.12	0.56	0.25
Panel C: Demographics						
Age	45.72	14.76	46.00	47.97	14.88	49.00
Female	0.52	0.50	1.00	0.50	0.50	1.00
White	0.64	0.48	1.00	0.67	0.47	1.00
Employed	0.78	0.42	1.00	0.78	0.41	1.00
Income	62555.00	44897.99	57500.00	61245.83	44057.29	42500.00
Financial Literacy	0.38	0.49	0.00	0.40	0.49	0.00
Observations	2,000			1,200		

Notes: Panel A reports summary statistics for survey-based expectations of economic outcomes, including expected stock returns, national and self-reported home price changes, inflation, mortgage rates, Treasury bill returns, wage growth, the probability of job loss, the probability of finding a job, and the probability of living to age 75. Panel B reports respondents beliefs about the relationships between stock returns and other asset returns. Panel C reports demographic characteristics of respondents, including age, gender, race, employment status, household income, and share of individuals that answered all four financial literacy questions correctly. The first three columns reports these statistics for all respondents and the last three columns reports them for individuals that responded in both the first and second round of surveys.

interquartile range spans from a "50-50" chance at the 25th percentile to an 83% probability at the 75th percentile. We observe similar degrees of disagreement regarding market outcomes. The standard deviation for expected stock market returns is 7.01 percentage points, more than half the average expected return of 12.31%. This pervasive disagreement across different eco-

nomic domains suggests that individuals do not merely differ in their specific outlooks for a single variable, but may hold fundamentally different internal models of the economy.

Comparing the full set of respondents to those who completed both rounds we find very similar expectations, beliefs, and demographics. Both groups believe stocks are positively correlated with t-bill rates, house prices, and wages but negatively correlated with inflation. Repeat respondents are slightly older, slightly more financially literate and make about $1,000 less per-year compared to the full set of respondents.

3.4 Correlation Structure

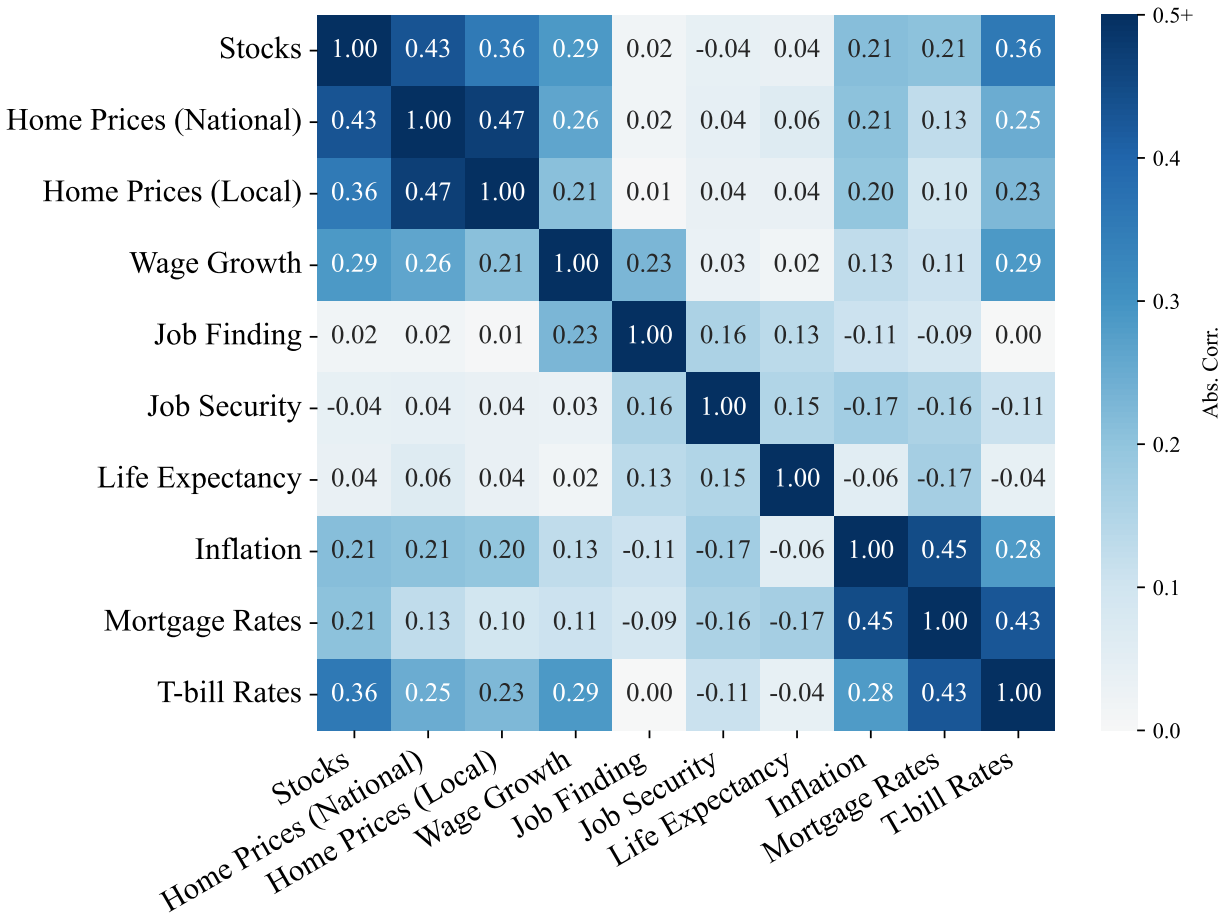
Figure 1 plots a heat map of pairwise correlations for the ten expectations elicited in our survey, with darker cells indicating larger absolute correlations. Three groups of beliefs co-move more within than across groups: assets (stocks and housing), rates (inflation, mortgage, and T-bill rates), and, to a lesser extent, human capital (job security, life expectancy, and wages). For example, stocks and national home prices have a correlation of 0.43, whereas stocks and job security (i.e., the probability of not losing one's job over the next year) are virtually uncorrelated at -0.04.

4 Three Facts about Correlated Beliefs

Fact 1: Beliefs co-move along a small number of common factors. We document that correlations across the ten elicited expectations are well-summarized by just three common factors, which together account for 53% of the total variation in beliefs. These three factors have an intuitive economic interpretation: an *asset factor* (loading on stock returns and house price growth), a *rates factor* (loading on interest rates and inflation), and a *human capital factor* (loading on wages, employment prospects, and life expectancy).

To extract this latent structure, we follow Stango and Zinman (2023) and conduct a common factor analysis rather than a principal component analysis. The key distinction is that the

Figure 1: Correlations of Expectations

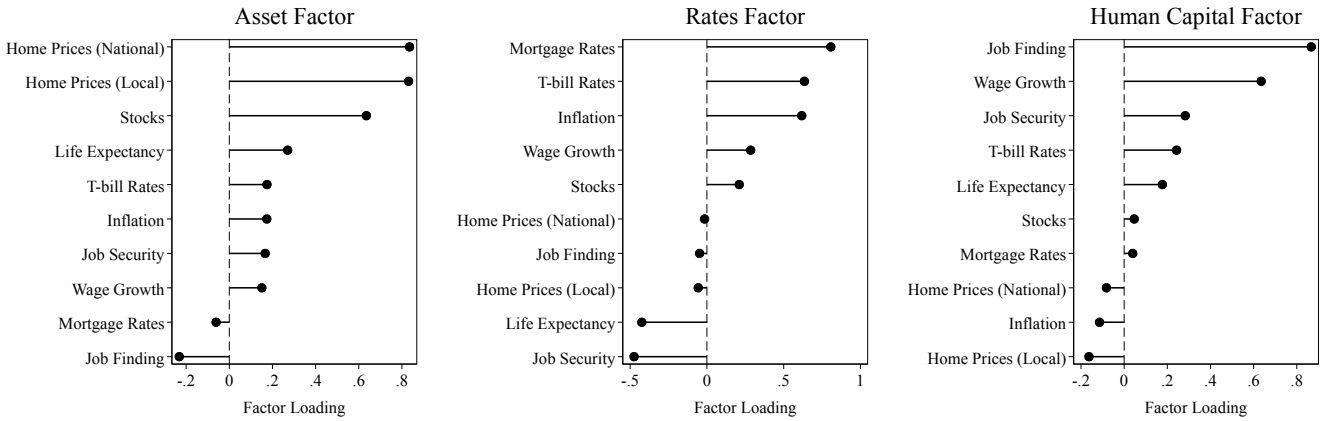


Notes: This heatmap shows pairwise correlations between our ten expectation variables, with annotated values indicating the raw correlations. Colors represent the magnitude of absolute correlations with darker colors indicating stronger relationships.

common factor model partitions the variance of each expectation into a shared latent component and an idiosyncratic "uniqueness" component, thereby isolating co-movement in beliefs from individual noise. This is the appropriate choice for our purposes: we want to characterize belief structures that are common across individuals, not reconstruct the full variance of each expectation. Formally, we decompose the variance of a standardized expectation matrix X (comprising p expectations) into m shared common factors:

$$\text{Var}(X) = \Lambda\Lambda' + D_\psi, \tag{3}$$

Figure 2: Factor Loadings by Expectation



Notes: This figure reports estimated factor loadings from a principal factor analysis with promax rotation on our 10 expectation variables. Each panel corresponds to one factor. Within each panel, variables are ordered by their estimated loading. The dot indicates the point estimate of the loading and the horizontal line connects the loading to zero.

where Λ is a $p \times m$ matrix of factor loadings and D_ψ is a diagonal matrix of uniquenesses. The (i,k)-th element of Λ, denoted λ_{ik}, represents the loading of expectation i on latent factor k. We standardize all expectations prior to estimation so that factor extraction depends on the underlying correlation structure rather than differences in scale or volatility. Under this representation, the covariance between any two distinct expectations i and j is determined entirely by their shared exposure to the latent factors:

$$\text{Corr}(X_i, X_j) = \sum_{k=1}^{m} \lambda_{ik}\lambda_{jk}, \qquad \text{for } i \neq j. \tag{4}$$

We estimate the model using principal factor analysis, which iteratively identifies factors to maximize explained common variance. Appendix B.3 describes the estimation in detail. Individual-level factor scores are estimated using the regression method as linear combinations of the standardized survey responses, with weights determined by the loading matrix and the inverse of the observed correlation matrix.

Figure 2 displays the estimated factor loadings λ_{ik}. We select factors with eigenvalues greater than one, a standard criterion (Stango and Zinman, 2023), retaining three factors

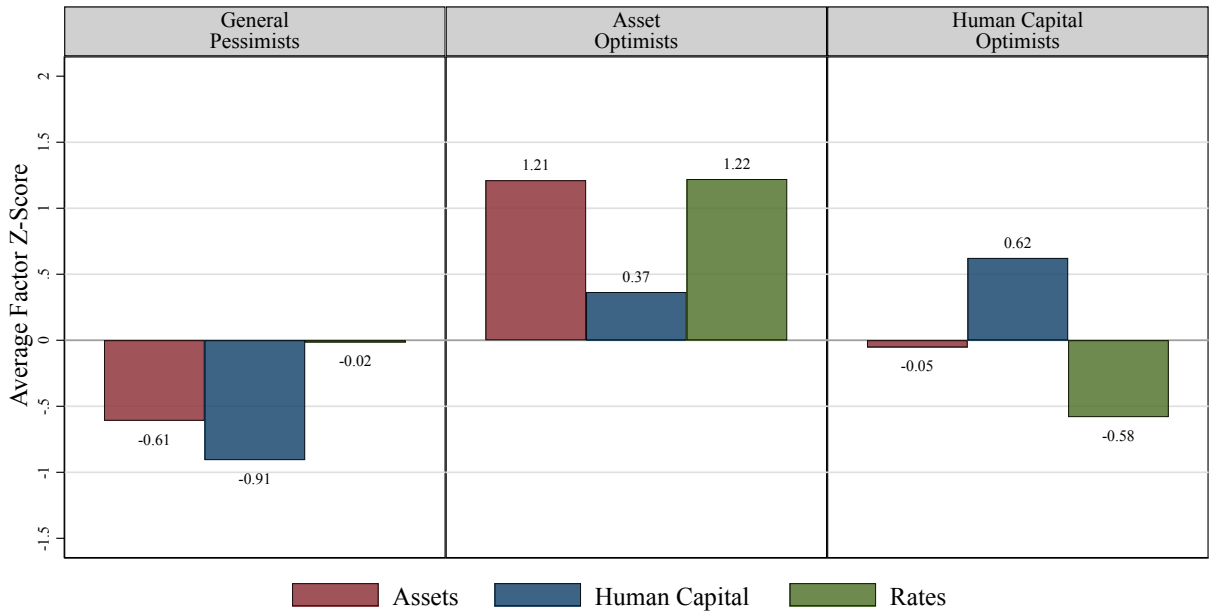
($m = 3$) with eigenvalues of 2.62, 1.43, and 1.08 that together explain 53% of the common variance. We label these three factors as follows: an *asset factor* (loading strongly on stock returns and home prices), a *rates factor* (loading on borrowing rates and inflation), and a *human capital factor* (loading on job prospects and wages).

The factor structure is economically intuitive. Individuals who believe stock returns will be high also tend to believe home prices will rise. Both are consistent with an optimistic view of future aggregate wealth. Beliefs about inflation and borrowing costs co-move because they share a common driver in expected monetary conditions. Human capital beliefs form a distinct cluster, consistent with individuals holding separate views about their own labor market prospects versus macro finance conditions. The degree to which beliefs cluster along these dimensions (rather than varying idiosyncratically) suggests that a low-dimensional characterization of belief structures is not merely a statistical convenience but reflects genuine latent heterogeneity in how individuals perceive the economic environment.

Fact 2: Belief correlations sort individuals into psychologically interpretable types. We next examine whether the common factors that structure beliefs also sort individuals into interpretable types. We identify three types that differ systematically in their belief profiles and correlate with independent measures of psychological optimism. Specifically, we perform a k-means clustering analysis over the three factor scores assigned to each individual. We determine the optimal number of clusters using the silhouette criterion, which measures the cohesion and separation of candidate clusterings. Figure 3 plots the average z-score for each factor score within each of the three clusters. Each factor captures a distinct dimension of beliefs, holding other factors constant: individuals with high asset factor scores believe stocks and housing will have high future returns; individuals with high human capital factor scores believe they will have high wage growth, job security, and life expectancy; individuals with high rate factor scores believe inflation and borrowing rates will be high in the future.

The three clusters correspond to economically interpretable belief types. The first group, *General Pessimists* (37%), expect low asset returns, weak human capital prospects, and moder-

Figure 3: K-means Clusters of Factor Scores



Notes: This figure plots the average standardized (z-score) factor score for our three factors by cluster. Clusters are identified using k-means clustering on individual-level factor scores. High asset score indicates higher expectations for asset returns. High human capital score indicates higher expectations for wage growth, life expectancy, and job security. High rates score indicates higher expectations for inflation and borrowing rates. Bars are color-coded by factor type. Demographic and expectations of each cluster, including age, income, financial literacy, and general optimism, are summarized in Table 2.

ate interest rates. The second group, *Asset Optimists* (21%), expects high asset returns, moderate human capital returns, and high interest rates. The third group, *Human Capital Optimists* (42%), expects strong labor market outcomes but moderate asset returns.

Table 2 reports demographics and average expectations for each belief type. There are meaningful differences across the groups. Optimists tend to earn more than pessimists, Human Capital Optimists are more financially literate than the other types, and pessimists are slightly older than the two optimist groups.

The most striking pattern in Table 2 concerns psychological optimism. Both optimists are at the 56th percentile of psychological optimism scores on average, while General Pessimists are at the 41st, a 15 percentage point gap. This alignment is notable because the psychological optimism measures were excluded entirely from both the factor analysis and the clustering

Table 2: Cluster Statistics

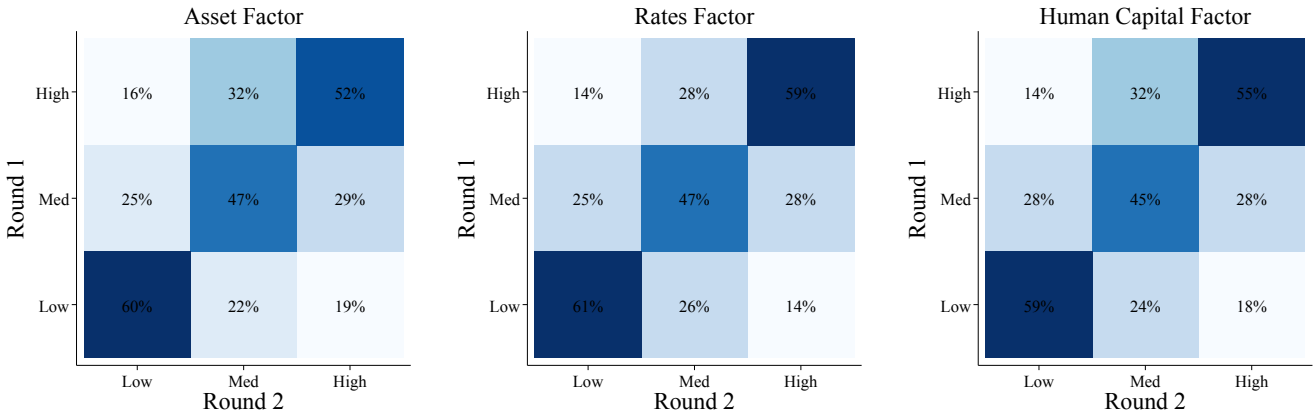
	General Pessimists	Asset Optimists	Human Capital Optimists
Panel A: Demographics			
Share of Sample	0.37	0.21	0.42
Mean Age	48.68	47.96	47.36
Mean Income	51185.68	60748.99	70375.49
Shr. Fin. Lit.	0.39	0.23	0.49
Optimism Perc.	41.11	55.83	56.33
Survey Length (Mins)	11.94	13.40	12.72
Shr. Meas. Error	0.41	0.52	0.33
Panel B: Avg. Expectations			
Stock returns	8.92	19.38	10.96
Home prices (national)	5.60	11.01	7.04
Home prices (local)	6.74	13.45	8.34
Wage growth	2.02	7.71	5.24
Find job Prob	27.49	47.65	69.46
Job Security Prob	68.26	72.23	86.75
Prob. Live to 75	58.02	64.04	74.49
Inflation	4.09	6.49	3.43
Mortgage Rates	7.56	11.20	6.31
T-bill Rates	4.23	9.36	4.21

Notes: This table reports demographic and expectation characteristics of the three clusters identified using k-means clustering on individual-level factor scores. Share financially literate is the share of individuals that answered all four financial literacy questions correctly. Optimism percentile is computed as the percentile of the average optimism score across the three psychological optimism questions. Share measurement error is the share of individuals that reported expected stock returns within one percentage point of their previous response to the same question.

procedure, so the correspondence between economic belief clusters and psychological disposition is a genuine out-of-sample validation of the cluster structure. In principle, psychological optimism should be orthogonal to expectations about macro outcomes such as stock returns and home prices. The fact that it is not suggests that individual psychology is a meaningful driver of the heterogeneous belief structures we document, and not merely an artifact of the statistical procedure used to recover them.

Fact 3: Beliefs are highly persistent within individuals, especially when correcting for

Figure 4: Factor Transitions

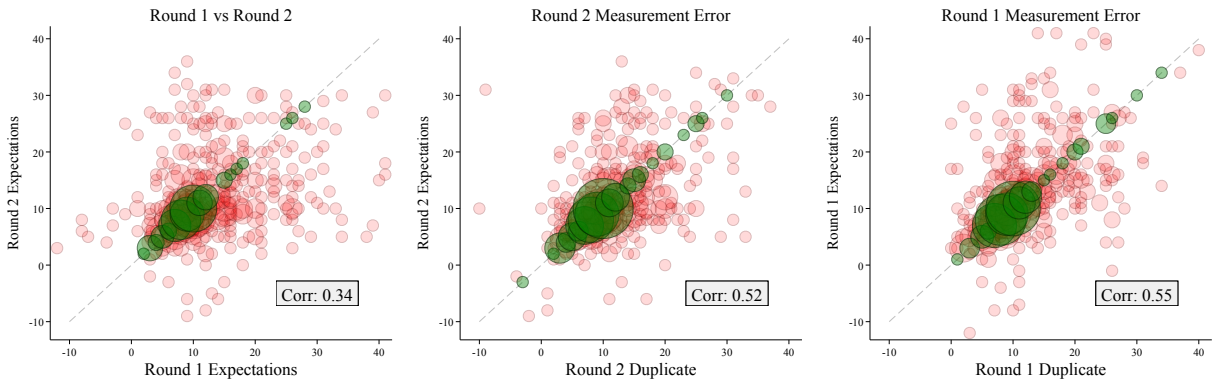


Notes: This figure shows the transitions of individual factor scores across survey rounds. Factor scores are partitioned into tertiles (Low, Medium, High) based on the full sample distribution. Each cell reports the percentage of respondents in a given Round 2 tertile (x-axis), conditional on their Round 1 tertile (y-axis). Diagonal elements represent the share of respondents with stable expectations across rounds. The sample is restricted to respondents who completed both survey rounds.

measurement error). We next ask how much beliefs vary within individuals over time. Our panel structure allows us to observe the same individuals across two survey rounds separated by six months. To characterize persistence, we divide individual factor scores into terciles for each round and plot the resulting transition matrix in Figure 4. Beliefs are remarkably stable. Among individuals in the lowest tercile of asset factor scores in Round 1, 60% remain in the lowest tercile in Round 2; persistence rates for the middle and top terciles are 47% and 52% respectively. We observe similar patterns for the rates and human capital factors. Transitions that do occur follow a monotone structure: individuals who move between terciles almost always move to an adjacent one, with long-range transitions being rare.

Raw cross-round correlations appear more moderate, but understate true stability: if beliefs are measured with error, genuine persistence will be attenuated toward zero. We exploit a re-elicitation design to separate measurement noise from true belief updating, eliciting stock return expectations twice within each survey session, once at the beginning and once at the end. Since no feedback occurs between the two elicitations, within-session disagreement likely

Figure 5: Changes in Stock Price Expectations



Notes: This figure plots stock return expectations across survey rounds and within two instances during the same survey. Panel A compares each respondent's expectation in Round 1 (July 2025) against their expectation in Round 2 (January 2026). Panels B and C compare two elicitations of the same question asked at the beginning and end of a single survey session, for Rounds 2 and 1 respectively. Because no feedback occurs between the two within-session elicitations, disagreement between them reflects response noise rather than genuine belief revision. Green bubbles indicate pairs of responses within 1 percentage point of each other; bubble size is proportional to the number of observations at that coordinate. Correlation coefficients are reported in the bottom-right corner of each panel.

reflects response noise rather than genuine belief revision. Figure 5 visualizes this. Panel A plots Round 1 against Round 2 stock return expectations. Panels B and C plot the initial elicitation against the re-elicitation for Rounds 2 and 1, respectively, with responses within one percentage point of each other shown in green.[2]

To quantify how much of the observed variation in belief changes can be attributed to noise, we model observed beliefs as

$$x_{ir} = \mu_{ir} + \epsilon_{ir} \tag{5}$$

where x_{ir} is the reported belief of individual i in round r, μ_{ir} is their true underlying belief, and ϵ_{ir} is mean-zero measurement error, independent of μ_{ir} and across rounds. The observed

[2]The within-survey correlation between the initial elicitation and the re-elicitation is approximately 0.50 in both rounds, well below one but not clustering near it, indicating that respondents are not simply anchoring their second response to their first. This validates the classical measurement error assumption we maintain throughout.

change in beliefs then decomposes as

$$\Delta x_i = (\mu_{i2} - \mu_{i1}) + (\epsilon_{i2} - \epsilon_{i1}) \tag{6}$$

so that

$$\text{Var}(\Delta x_i) = \text{Var}(\Delta \mu_i) + \sigma^2_{\epsilon,1} + \sigma^2_{\epsilon,2}. \tag{7}$$

The re-elicitations identify $\sigma^2_{\epsilon,r}$ without any cross-round assumption. Letting x'_{ir} denote the second within-survey elicitation, we have

$$x_{ir} - x'_{ir} = \epsilon_{ir} - \epsilon'_{ir} \implies \text{Var}(x_{ir} - x'_{ir}) = 2\sigma^2_{\epsilon,r} \tag{8}$$

so that $\hat{\sigma}^2_{\epsilon,r} = \frac{1}{2}\text{Var}(x_{ir} - x'_{ir})$, estimated separately for each round.

We estimate the measurement error variance to be 22.46 in Round 1 and 21.20 in Round 2. The near-identical values across rounds indicate that the survey instrument is stable and that respondents are no noisier in one round than the other. The observed variance of cross-round changes in stock return expectations is 69.34, implying that 63% of the variation in observed belief changes is attributable to measurement error rather than genuine updating:

$$\hat{\pi}_{ME} = \frac{\hat{\sigma}^2_{\epsilon,1} + \hat{\sigma}^2_{\epsilon,2}}{\text{Var}(\Delta x_i)} = \frac{22.46 + 21.20}{69.34} = 0.630. \tag{9}$$

Removing this noise, the denoised cross-round correlation rises from 0.34 to 0.58, consistent with the high persistence documented in Figure 4.[3]

[3] Specifically, $\hat{\rho}_{\text{true}} = \text{Cov}(x_{i1}, x_{i2}) / \sqrt{\text{Var}(\hat{\mu}_{i1})\text{Var}(\hat{\mu}_{i2})}$, where $\text{Cov}(x_{i1}, x_{i2}) = \text{Cov}(\mu_{i1}, \mu_{i2})$ under independence of measurement errors across rounds.

5 Life Cycle Model

We build and simulate a rich life cycle portfolio choice model disciplined by the belief structures documented in Section 4. The model extends the consumption-saving framework of Choukhmane and de Silva (2024) and Choukhmane (2025) to incorporate heterogeneous beliefs over asset returns, labor market outcomes, life expectancy, and the correlation structure of returns. Agents choose consumption, savings, portfolio allocations between a risk-free bond and a risky asset, as well as housing tenure. The key departure from a standard life cycle model is that agents solve their problem under subjective beliefs that may deviate from the true data-generating process in all of these dimensions simultaneously.

The remainder of this section proceeds as follows. Sections 5.1–5.5 describe the economic environment. Section 5.6 defines the belief system agents use to make decisions. Section 5.7 states the household's dynamic optimization problem. Section 5.8 describes calibration, solution, and simulation.

5.1 Demographics and Preferences

Each period corresponds to one year. Agents enter working life at age a_0, work for T_w periods, retire, and live at most T periods total. Before their certain death at $t = T$, agents face age-dependent mortality risk with m_t denoting the probability of surviving to period $t+1$ conditional on being alive at t. We denote age as $a_t = t + a_0$.

Agents have CRRA preferences with a Cobb-Douglas aggregator over non-durable consumption C_t and housing services H_t, with housing preference weight ν. Time discounting is β and relative risk aversion is γ. Per-period utility is adjusted by an equivalence scale n_t that captures changes in household size over the life cycle (Lusardi et al., 2017)

$$U(C_t, H_t) = \frac{n_t \left(\frac{C_t^{1-\nu} H_t^{\nu}}{n_t} \right)^{1-\gamma}}{1-\gamma}. \tag{10}$$

5.2 Labor Market

At any point in time, an agent's employment status emp_t takes one of three values: employed (E), unemployed (U), or retired (Ret). All agents retire deterministically at $t = T_w$. During their working life, agents face uncertainty over employment status and wage realizations.

Employment ($emp_t = E$). Employed agents earn exogenous income w_t consisting of a deterministic component cubic in age and a persistent stochastic component:

$$\begin{aligned} \ln w_t &= \delta_0 + \delta_1 a_t + \delta_2 a_t^2 + \delta_3 a_t^3 + \eta_t, \\ \eta_t &= \rho\, \eta_{t-1} + \xi_t^E, \quad \xi_0^E \sim \mathcal{N}(0, \sigma^2_{\xi_0}), \quad \xi_t^E \sim \mathcal{N}(0, \sigma^2_{\xi}) \;\; \forall\, t > 0. \end{aligned} \tag{11}$$

Unemployment ($emp_t = U$). An employed agent becomes unemployed with age-dependent probability $\pi^U(t)$ and receives unemployment benefits $ui_t = ui(\eta_t)$ based on their most recent earnings state. Upon re-employment, income evolves as

$$\ln w_{t+1} = \delta_0 + \delta_1 a_{t+1} + \delta_2 a_{t+1}^2 + \delta_3 a_{t+1}^3 + \eta_{t+1}, \quad \eta_{t+1} = \rho\, \eta_t + \xi^U_{t+1}, \quad \xi^U_{t+1} \sim \mathcal{N}(\mu^{UE}, \sigma^2_{\xi}), \tag{12}$$

where $\mu^{UE} < 0$ captures the persistent wage scarring associated with unemployment spells. Unemployed agents find work in the next period with age-dependent probability $\pi^E(t)$.

Retirement ($emp_t = Ret$). All agents retire at $t = T_w$ and receive public pension benefits $ss_t = ss(ae)$, where ae denotes average lifetime earnings at retirement.

5.3 Government

Unemployment benefits. Upon separation, agents receive $ui(\eta_t)$, which depends on the persistent earnings component from their last period of employment.

Retirement benefits. Social security benefits $ss(ae)$ are a function of average lifetime earnings.

5.4 Financial Assets

There are three financial assets. A *risk-free bond* pays a constant gross return $R^B = R_f$ per year. A *risky asset* represents a diversified equity index with stochastic i.i.d. log returns:

$$\ln R_t^S = \ln R_f + \mu_s + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_s^2). \tag{13}$$

Finally, agents may purchase *housing* at price P_t^H, which grows at a constant gross rate R^H per year such that $P_{t+1}^H = R^H P_t^H$. In the true data-generating process, equity returns, housing returns, and labor income shocks are uncorrelated. As described in Section 5.6, agents may hold heterogeneous beliefs about these correlations.

5.5 Housing

An agent's housing status $house_t$ is either owner (O) or renter (R) in each period.

Renters ($house_t = R$). Renters pay annual rent

$$X_t(H_{it}) = \frac{P_t^H H_{it}}{\chi},$$

where H_{it} is the quantity of housing services chosen and χ is the price-to-rent ratio.

Owners ($house_t = O$). Renters may purchase a home financed by a fixed-rate mortgage at rate r^M with length L^M periods. Purchase requires a down payment of $\phi P_t^H H_{it}$, followed by fixed annual payments

$$A = \frac{(1-\phi) P_t^H H_{it}}{\frac{1}{r^M}\left(1 - \frac{1}{(1+r^M)^{L^M}}\right)}. \tag{14}$$

Owners may sell their home at any time for its current market value $P_t^H H_{it}$, subject to a proportional transaction cost $c_H P_t^H H_{it}$ and repayment of any remaining mortgage principal.[4] Beliefs about home price growth affect the rent-vs-buy decision and the perceived value of owning we describe this channel in Section 5.6.

[4] We calibrate the model so that agents can always afford to sell their home. Mortgage default is not permitted.

5.6 Belief Structure

Agents have heterogeneous beliefs about asset returns, labor market outcomes, survival probabilities, and the correlation structure of returns. This section defines the four layers of the belief system, from the true data-generating process through to the subjective parameters agents use to make decisions. How each layer is disciplined empirically is deferred to Section 5.8.1.

Layer 1: True Parameters

The true data-generating process is governed by the parameters described in Sections 5.1–5.5: average equity returns μ_s, risk-free rate R_f, housing return R_H, the wage process $\{\delta_k, \rho, \sigma_\xi\}$, mortality $\{m_t\}$, and employment transition probabilities $\{\pi^U(t), \pi^E(t)\}$. Critically, equity returns, labor income shocks, housing returns, and risk free rates are *mutually independent* in the data-generating process:

$$\text{Cov}(R_t^S, w_t) = \text{Cov}(R_t^S, R_t^H) = \text{Cov}(R_t^S, R_t^B) = 0.$$

Any perceived non-zero correlation between these quantities therefore represents a pure belief distortion with no rational hedging foundation.

Layer 2: Type-Level Beliefs

We partition agents into a finite set of belief types $\mathcal{J}$. Treating the population mean belief as correct, the belief distortion of type j over outcome $x \in \{s, b, \pi, h, w, m, U, E\}$ is defined as the average deviation of that type's members from the population mean,

$$\epsilon_j^x = \frac{1}{N_j} \sum_{i \in j} x_i - \bar{x}, \tag{15}$$

where N_j is the number of individuals in type j and $\bar{x}$ is the population mean belief over outcome x. We additionally define type-level perceived correlations $\rho_j^{s,w}$, $\rho_j^{s,h}$, and $\rho_j^{s,b}$, rep-

resenting the average perceived correlations between equity returns and labor income, home prices, and the risk-free rate respectively within each type. How types are constructed and how these moments are estimated from the survey is described in Section 5.8.1.

Layer 3: Individual Draws Around the Type Mean

Within each type, agents draw individual-specific belief realizations:

$$\epsilon^x_{i,j,t} = \epsilon^x_j + \eta^x_{i,j,t}, \quad \eta^x_{i,j,t} \sim \mathcal{N}(0, \sigma^x_j), \tag{16}$$

where σ^x_j is the within-type dispersion of beliefs over outcome x. Equation 16 decomposes individual beliefs into two components: a persistent type-level mean ε^x_j and a transitory idiosyncratic deviation $\eta^x_{i,j,t}$. This decomposition has a natural implication for the dynamic program. Because the idiosyncratic component is transitory, it shifts an agent's current beliefs without affecting their expectations about future beliefs. Only the persistent component ε^x_j therefore enters the value function as a state variable, keeping the state space tractable while still allowing for belief heterogeneity within types.

Layer 4: Perceived Parameters Used in Decision-Making

Agents make current decisions using their individual belief draws $\epsilon^x_{i,j,t}$, but continuation values depend only on the persistent type j, as established in Layer 3. For asset returns and wages, perceived nominal beliefs are adjusted for perceived inflation, so that decisions depend on perceived *real* returns:

$$\begin{aligned}
\tilde{\mu}_{s,ij} &= \mu_s + (\epsilon^s_{i,j,t} - \epsilon^\pi_{i,j,t}) & \tilde{m}_{t,ij} &= m_t \cdot (1 + \epsilon^m_{i,j,t}) \\
\tilde{R}^H_{ij} &= R^H + (\epsilon^h_{i,j,t} - \epsilon^\pi_{i,j,t}) & \tilde{\pi}^U_{ij}(t) &= \pi^U(t) + \epsilon^U_{i,j,t} \\
\tilde{R}^B_{ij} &= R_f + (\epsilon^b_{i,j,t} - \epsilon^\pi_{i,j,t}) & \tilde{\pi}^E_{ij}(t) &= \pi^E(t) + \epsilon^E_{i,j,t} \\
\Delta \ln \tilde{w}_{ij} &= \Delta \ln w_t + (\epsilon^w_{i,j,t} - \epsilon^\pi_{i,j,t})
\end{aligned} \tag{17}$$

In addition, agents perceive the covariance matrix of returns to have off-diagonal elements $\tilde{\rho}_j^{s,w}$, $\tilde{\rho}_j^{s,h}$, and $\tilde{\rho}_j^{s,b}$, which affect portfolio allocation even though the true covariances are zero. When making decisions, agents take their current belief type as fixed and the full belief set for agent i of type j at time t is denoted:

$$\text{belief}_{i,j,t} = \left\{ \epsilon^s_{i,j,t},\ \epsilon^b_{i,j,t},\ \epsilon^\pi_{i,j,t},\ \epsilon^h_{i,j,t},\ \epsilon^w_{i,j,t},\ \epsilon^m_{i,j,t},\ \epsilon^U_{i,j,t},\ \epsilon^E_{i,j,t},\ \rho_j^{s,w},\ \rho_j^{s,h},\ \rho_j^{s,b} \right\}. \tag{18}$$

5.7 Individual Problem

State Space

The state vector for agent i of type j at time t is:

$$\mathbf{s}_{it} = \left\{ \underbrace{a_t}_{\text{age}},\ \underbrace{W_{it}}_{\text{liquid wealth}},\ \underbrace{\eta_{it}}_{\text{earnings state}},\ \underbrace{emp_t}_{\text{employment}},\ \underbrace{house_t}_{\text{housing status}},\ \underbrace{H_{it}}_{\text{housing quantity}},\ \underbrace{M_{it}}_{\text{mortgage principal}},\ \underbrace{j_t}_{\text{belief type}} \right\}. \tag{19}$$

Age a_t and belief type j_t are discrete. Liquid wealth W_{it}, the persistent earnings component η_{it}, housing quantity H_{it}, and remaining mortgage principal M_{it} are continuous. Employment status $emp_t \in \{E, U, Ret\}$ and housing status $house_t \in \{O, R\}$ are discrete.

Choice Variables

In each period, agents choose: (i) non-durable consumption C_{it}; (ii) end-of-period liquid wealth W_{it+1}; (iii) portfolio share allocated to equities $\alpha_{it} \in [0, 1]$; (iv) housing tenure decision $\in$ {stay, buy, sell}; and (v) housing quantity H_{it+1} for renters and new buyers.

Decision Problem

Because agents are naive about future belief evolution, the value function is solved separately for each belief type j, taking belief$_j$ as fixed and time-invariant. For an employed agent ($emp_t =$

E):

$$V_j(\mathbf{s}_{it}) = \max_{C_{it}, W_{it+1}, \alpha_{it}, house_{t+1}, H_{it+1}} U(C_{it}, H_{it}) + \beta\, \tilde{m}_{t,j}\, \mathbb{E}_j\left[V_j(\mathbf{s}_{it+1}) \,\middle|\, \mathbf{s}_{it},\ \text{belief}_j\right], \tag{20}$$

where $\mathbb{E}_j[\cdot]$ denotes expectations formed under the subjective belief set of type j and $\tilde{m}_{t,j}$ is the perceived survival probability. The perceived portfolio return between t and $t+1$ is

$$\tilde{R}^P_{it+1} = \alpha_{it}\tilde{R}^S_{t+1} + (1-\alpha_{it})\tilde{R}^B_j, \tag{21}$$

where the agent perceives equity and bond returns to covary with coefficient $\tilde{\rho}^{s,b}_j$ and perceives both to covary with income shocks and housing returns according to $\tilde{\rho}^{s,w}_j$ and $\tilde{\rho}^{s,h}_j$ respectively. The budget constraint (when the agent doesn't change their housing status) is:

$$W_{it+1} = \tilde{R}^P_{it+1}\left(W_{it} + w_t - C_{it} - X_t(H_{it}) \cdot \mathbf{1}[house_t = R] - A_{it} \cdot \mathbf{1}[house_t = O]\right), \tag{22}$$

where $X_t(H_{it})$ is rent paid by renters and A_{it} is the mortgage payment made by owners. The decision problem for the unemployed and retired cases are analogous, replacing w_t with unemployment benefits $ui(\eta_t)$ and social security $ss(ae)$ respectively.

5.8 Bringing the Model to the Data

5.8.1 Calibration

We calibrate our model using common parameter choices from the literature. Table A1 reports the parameter values and sources for each of the key parameters in the model. Each is discussed briefly below.

Demographics. We set the life cycle length, retirement age, and maximum lifespan to match standard targets in the household finance literature. Mortality risk is calibrated to U.S.

actuarial data. The equivalence scale is estimated to capture changes in household composition over the life cycle.

Labor income process. We estimate the parameters of the income process using data from the Survey of Income and Program Participation (SIPP) at the annual frequency. The deterministic and stochastic components of equation (11) are estimated jointly using a minimum distance approach with measurement error correction (following Guvenen, 2009). Estimates are consistent with prior literature, with a relatively high estimated persistence of permanent income shocks. Employment transition probabilities are estimated from SIPP microdata.

Asset returns. We calibrate asset return parameters to long-run U.S. historical averages. In the data-generating process, asset returns are uncorrelated with income shocks and with each other. Perceived correlations are entirely a product of the belief distortions elicited from the survey.

Beliefs. We construct belief types by discretizing individual factor scores from Section 4 into terciles and defining types as all combinations of tercile assignments across the three factors, yielding 27 types. Type-level belief distortions ϵ_j^x are defined as deviations from the sample mean belief, which we treat as correct. Within-type belief dispersion σ_j^x is set to the empirical within-type standard deviation for each outcome and type. Perceived correlation parameters $\rho_j^{s,w}$, $\rho_j^{s,h}$, $\rho_j^{s,b}$ are set to the average perceived correlation among individuals in each type, elicited directly from the survey.

5.8.2 Solving the Model

We solve the model by value function iteration separately for each of the 27 belief types. For type j, the value function in equation (20) is solved taking belief_j as fixed truth. The state space is discretized using standard methods.

5.8.3 Simulation

We simulate $N = 5{,}000$ individuals. Each agent is assigned an initial belief type j at $t = 0$ according to the empirical distribution of types in Round 1 of our survey. In each subsequent period, belief type evolves according to the empirical transition matrix estimated from the Round 1 to Round 2 panel transitions.

In each period, agent i of type j draws an individual belief realization $\epsilon^{x}_{i,j,t}$ according to equation (16), observes their employment and income states, and makes optimal consumption, portfolio, and housing decisions under belief$_j$. Liquid wealth, housing equity, and the mortgage balance are updated accordingly.

5.8.4 Validation

The key moments we aim to match are the correlation structure of beliefs found in the survey data. To test how well we match these, Figure A2 plots the belief type level correlations of the expectations found in the data (left) next to those in our simulated data (right). The matrix is, by construction, very similar with many of the moments matching exactly.

6 The Behavioral and Welfare Implications of Correlated Beliefs

6.1 Correlations Dampen the Pass-Through from Beliefs to Behavior

Correlated beliefs introduce a classic omitted variable bias when estimating the effect of each one belief on behavior. The direction and magnitude of the resulting bias depends on the sign and strength of belief correlations. In our setting, the observed correlation structure generally implies attenuation, though not uniformly across all outcomes.

To isolate this effect, we run two sets of simulations. In the first, we introduce belief heterogeneity one domain at a time, holding all others beliefs homogeneous. In the second, all beliefs

vary jointly according to their empirical distribution. Comparing regressions of outcomes on beliefs across these two calibrations quantifies the extent to which correlated beliefs attenuate or amplify the pass-through from any single belief to behavior.

Table 3 presents the results. Columns (1) and (2) report coefficients from the isolated-belief and full-belief simulations respectively, and column (3) reports their ratio. A ratio below one indicates attenuation while a ratio above one indicates amplification. The pattern across outcomes is broadly consistent with attenuation. A one percentage point increase in expected stock returns raises the equity share by 3.09 percentage point when all other beliefs are held homogeneous. This pass-through coefficient falls to 2.45 percentage points when the full correlation structure is active, a reduction of 21%. The attenuation is even more pronounced for the pass-through from beliefs about bond returns to equity share: the coefficient moves from -4.45 to -0.58, an 87% decline in absolute value. For consumption, the pass-through from life expectancy and job loss beliefs is similarly attenuated, while job finding and wage growth beliefs exhibit a ratio above one, indicating amplification.

6.2 Correlations Reduce the Welfare Impact of Belief Heterogeneity

When beliefs are correlated across domains, their distortions tend to offset one another, reducing the welfare cost of belief heterogeneity. Holding fixed the dispersion of each belief, the empirical correlation structure reduces average lifetime welfare losses by nearly one-third relative to a benchmark with uncorrelated beliefs.

To isolate the role of correlation, we compare the baseline calibration, in which beliefs are correlated across domains as in the data, to a counterfactual with the same mean and standard deviation of each belief but no cross-domain correlation. In both cases, the population is correct on average, so the exercise isolates the welfare consequences of the correlation structure rather than the level of beliefs. We keep fixed the 27-type structure, type frequencies, and transition probabilities. In the counterfactual (*Random Beliefs*), each type's belief vector is drawn independently across domains using a Gram-Schmidt orthogonalization that removes

Table 3: Belief Pass-through to Household Decisions

		All Belief Types		
		(1)	(2)	(3)
Outcome	Belief Type	Isolated	All	All ÷ Isolated
Panel A: Portfolio Allocation				
Change in Equity Share	Stock Return Belief	3.09	2.45	0.79
	Bond Return Belief	-4.45	-0.58	0.13
Panel B: Housing				
Change in Homeownership (×100)	House Price Belief	0.69	0.11	0.16
Panel C: Consumption				
Change in Log Consumption	Wage Growth Belief	2.19	2.46	1.12
	Job Loss Belief	-0.38	-0.32	0.86
	Job Finding Belief	0.20	0.29	1.47
	Survival Belief	-1.64	-0.90	0.55

Notes: Columns (1) and (2) report OLS coefficients from regressions of the outcome on the stated belief measure. Column (1) uses simulated data in which all beliefs are homogeneous except the belief used as the explanatory variable; column (2) uses simulated data in which all beliefs are calibrated to match the heterogeneity and correlations in the survey responses. Column (3) reports the ratio of column (2) to column (1); values below one indicate attenuation of pass-through when the full belief structure is active. Homeownership coefficients are multiplied by 100. Wage growth and job loss regressions are restricted to employed individuals under age 65. Job finding regressions are restricted to unemployed individuals under age 65.

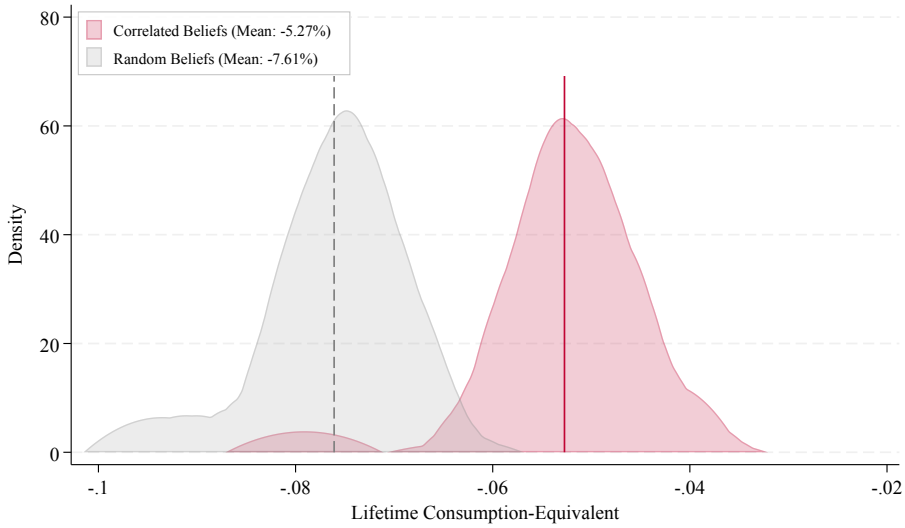
cross-belief correlations while preserving each belief's marginal moments exactly, and agents re-solve the full dynamic problem under the new belief system.

For each type, we measure welfare as the lifetime consumption-equivalent variation relative to a benchmark agent with fully correct and homogeneous beliefs,

$$CEV_k = \left(\frac{V_k}{V^*}\right)^{\frac{1}{1-\gamma}} - 1,$$

where V_k denotes lifetime utility for type k and V^* denotes lifetime utility under homogeneous correct beliefs. Lifetime utility is constructed from simulated paths of non-housing consumption and housing services using the model's period utility aggregator and then averaged within

Figure 6: Belief Correlations Attenuate Welfare Losses



Notes: This figure plots the distribution of lifetime consumption-equivalent welfare losses across the 27 simulated belief types, relative to a benchmark with homogeneous correct beliefs. In the *Correlated Beliefs* calibration, beliefs match the empirical joint distribution of survey responses. In the *Random Beliefs* calibration, beliefs are drawn independently across domains using a Gram-Schmidt orthogonalization that removes cross-belief correlations while preserving each belief's mean and dispersion. In both environments, the 27-type structure, type frequencies, and transition probabilities are held fixed, and agents re-solve the full dynamic problem. Vertical lines indicate the mean of each distribution. Values further to the left indicate larger welfare losses.

type. Negative values indicate welfare losses from belief heterogeneity relative to the homogeneous benchmark. Figure 6 plots the distribution of lifetime welfare losses under both calibrations. Mean welfare losses are −5.27% under correlated beliefs, compared with −7.61% under random beliefs, a reduction of 2.34 percentage points, or 30.7% relative to the random-beliefs benchmark. The attenuation is broad-based across the entire distribution. The correlated-beliefs distribution shifts uniformly to the right: the 10th percentile improves from −7.41% to −5.22%, the median from −6.66% to −4.46%, and the 90th percentile from −6.04% to −3.80%. The approximately uniform gap of 2.2 percentage points across percentiles indicates that the result is not driven by a small number of extreme types but reflects a broad shift in the welfare distribution.

The intuition parallels the pass-through attenuation results in Section 6.1. When beliefs are correlated, optimism in one domain tends to co-occur with optimism in others, and the re-

sulting behavioral distortions partially offset one another. As the simple framework in Section 2 shows, income optimism and longevity optimism push consumption in opposite directions, so a household that is optimistic about both may behave close to optimally even though each belief is individually distorted. Randomly recombining equally dispersed beliefs removes these offsets and allows distortions to reinforce one another, substantially amplifying welfare losses. Studying beliefs one domain at a time implicitly imposes this random-recombination assumption and therefore overstates the welfare cost of belief heterogeneity.

7 Conclusion

This paper documents the joint distribution of household subjective expectations and its implications for lifecycle financial behavior. Incorporating the empirical belief correlations into a lifecycle model delivers two main results. First, the observed correlation structure substantially attenuates the pass-through from any single expectation to behavior. Correlated beliefs thus contribute to explaining the puzzling lack of sensitivity of behavior to self-reported expectations. Second, it reduces the welfare cost of belief heterogeneity by nearly one-third relative to a benchmark with equally dispersed but uncorrelated beliefs. Studying expectations one domain at a time overstates both the behavioral and welfare consequences of belief heterogeneity.

An open question is why belief correlations take the form we document. One possibility is that belief formation is shaped by learning and adaptation. If some correlations of beliefs generate especially costly mistakes over the life cycle, those joint distortions may be more likely to be corrected over time through personal experience, social learning, or the transmission of adaptive heuristics. On this view, the attenuating correlation structure we measure may partly reflect a long-run equilibrium in which the most harmful combinations of belief errors are gradually corrected.

References

Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D Shapiro, and Christopher Tonetti, "Long-term-care utility and late-in-life saving," *Journal of Political Economy*, 2020, *128* (6), 2375–2451.

Armona, Luis, Andreas Fuster, and Basit Zafar, "Home price expectations and behaviour: Evidence from a randomized information experiment," *The Review of Economic Studies*, 2019, *86* (4), 1371–1410.

Bachmann, Ruediger, Giorgio Topa, and Wilbert van der Klaauw, *Handbook of economic expectations*, Elsevier, 2022.

Bergman, Abigail, Alex Chinco, Samuel M Hartzmark, and Abigail B Sussman, "Survey curious? start-up guide and best practices for running surveys and experiments online," *Start-Up Guide and Best Practices For Running Surveys and Experiments Online (October 05, 2020)*, 2020.

Caplin, Andrew, "Data engineering for cognitive economics," *Journal of Economic Literature*, 2025, *63* (1), 164–196.

_ , Victoria Gregory, Eungik Lee, Søren Leth-Petersen, and Johan Sæverud, "Subjective earnings risk," Technical Report, National Bureau of Economic Research 2023.

Carranza, Guillermo, Taha Choukhmane, Fiona Greig, Cormac O'Dea, and Lawrence Schmidt, "Surveying Counterfactuals: Improving 401(k) Matches Using Hypothetical Choices," 2026.

Catherine, Sylvain, "Countercyclical labor income risk and portfolio choices over the life cycle," *The Review of Financial Studies*, 2022, *35* (9), 4016–4054.

Chapman, Jonathan, Mark Dean, Pietro Ortoleva, Erik Snowberg, and Colin Camerer, "Econographics," *Journal of Political Economy Microeconomics*, 2023, *1* (1), 115–161.

Choukhmane, Taha, "Default options and retirement saving dynamics," *American Economic Review*, 2025, *115* (11), 3749–3787.

_ and Tim de Silva, "What Drives Investors' Portfolio Choices? Separating Risk Preferences from Frictions," Technical Report, National Bureau of Economic Research 2024.

D'Acunto, Francesco, Evangelos Charalambakis, Dimitris Georgarakos, Geoff Kenny, Justus Meyer, and Michael Weber, "Household inflation expectations: An overview of recent insights for monetary policy," 2024.

Enke, Benjamin, Thomas Graeber, Ryan Oprea, and Jeffrey Yang, "Behavioral attenuation," Technical Report, National Bureau of Economic Research 2024.

Fabrigar, Leandre R and Duane T Wegener, *Exploratory factor analysis*, Oxford University Press, 2012.

Fuster, Andreas and Basit Zafar, "Survey experiments on economic expectations," in "Handbook of Economic Expectations," Elsevier, 2023, pp. 107–130.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, "Five facts about beliefs and portfolios," *American Economic Review*, 2021, *111* (5), 1481–1522.

Gorodnichenko, Yuriy and Xiao Yin, "Higher-order beliefs and risky asset holdings," Technical Report, National Bureau of Economic Research 2024.

Guvenen, Fatih, "An empirical investigation of labor income processes," *Review of Economic dynamics*, 2009, *12* (1), 58–79.

Hartzmark, Samuel M and Abigail B Sussman, "Eliciting Expectations," *Available at SSRN*, 2024.

Heimer, Rawley Z, Kristian Ove R Myrseth, and Raphael S Schoenle, "YOLO: Mortality beliefs and household finance puzzles," *The Journal of Finance*, 2019, *74* (6), 2957–2996.

Hurd, Michael D and Kathleen McGarry, "Evaluation of the subjective probabilities of survival in the health and retirement study," *Journal of Human resources*, 1995, pp. S268–S292.

Indarte, Sasha, Raymond Kluender, Ulrike Malmendier, and Michael Stepner, "Consumption Wedges: Measuring and Diagnosing Distortions," 2025.

Koşar, Gizem and Cormac O'Dea, "Expectations data in structural microeconomic models," in "Handbook of Economic Expectations," Elsevier, 2023, pp. 647–675.

Liu, Haoyang and Christopher Palmer, "Implicit Extrapolation and the Beliefs Channel of Investment Demand," Technical Report, Working Paper 2023.

_ and _, "Implicit extrapolation and the beliefs channel of investment demand," *Journal of Financial Economics*, 2026, *175*, 104203.

Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S Mitchell, "Optimal financial knowledge and wealth inequality," *Journal of political Economy*, 2017, *125* (2), 431–477.

Mueller, Andreas I, Johannes Spinnewijn, and Giorgio Topa, "Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias," *American Economic Review*, 2021, *111* (1), 324–363.

O'Dea, Cormac and David Sturrock, "Survival pessimism and the demand for annuities," *Review of Economics and Statistics*, 2023, *105* (2), 442–457.

Palan, Stefan and Christian Schitter, "Prolific. ac—A subject pool for online experiments," *Journal of behavioral and experimental finance*, 2018, *17*, 22–27.

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

Scheier, Michael F and Charles S Carver, "Optimism, coping, and health: assessment and implications of generalized outcome expectancies.," *Health psychology*, 1985, *4* (3), 219.

Stango, Victor and Jonathan Zinman, "We are all behavioural, more, or less: A taxonomy of consumer decision-making," *The Review of Economic Studies*, 2023, *90* (3), 1470–1498.

Stantcheva, Stefanie, "How to run surveys: A guide to creating your own identifying variation and revealing the invisible," *Annual Review of Economics*, 2023, *15* (1), 205–234.

Velásquez-Giraldo, Mateo, "Life-Cycle Portfolio Choices and Heterogeneous Stock Market Expectations," 2024.

A Additional Tables and Figures

Table A1: Calibrated Parameters

	Parameter	Value	Description
Preferences			
	γ	2.50	Coefficient of risk aversion
	ν	0.30	Housing preference weight
Financial Market			
	μ_s	0.08	Average equity return
	σ_s	0.25	Equity return s.d.
	R_f	0.02	Risk-free rate
Housing Market			
	R^H	0.03	Housing return
	χ	22.00	Price-to-rent ratio
	ϕ	0.10	Down payment
	c_H	0.10	Transaction cost
Demographics			
	a_0	21.00	First age
	T_w	44.00	Working years
	T	79.00	Total periods
Income Process			
	δ_0	1.63	
	δ_1	0.09	
	δ_2	-1.07e-3	
	δ_3	3.90e-7	
	ρ	0.98	
	σ_ξ	0.11	

Notes: This table reports the parameter values used in the baseline calibration. Housing parameters and preferences follow Catherine (2022). Income profile estimates are based on SIPP data as reported in Choukhmane and de Silva (2024).

Figure A1: Survey Instrument

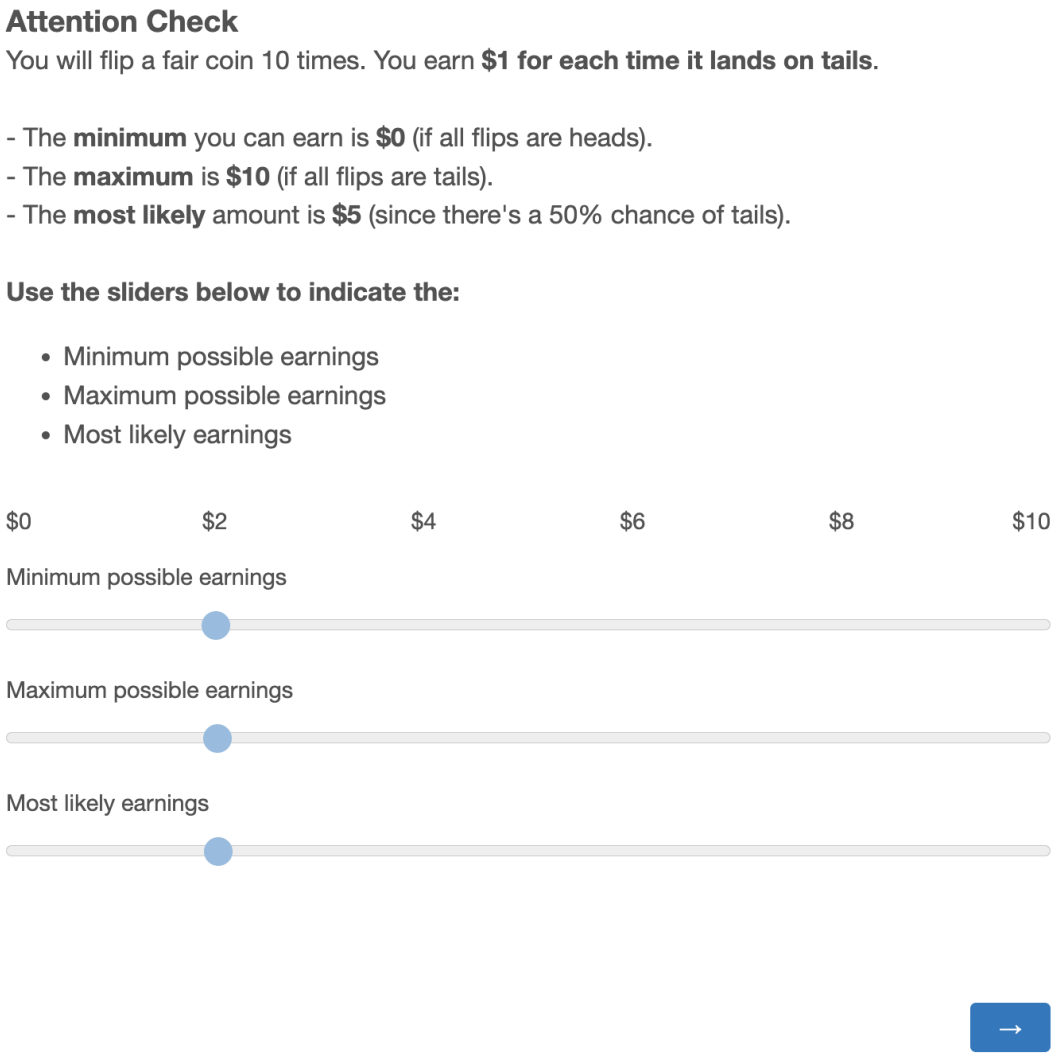


Figure A1: Survey Instrument (Continued)

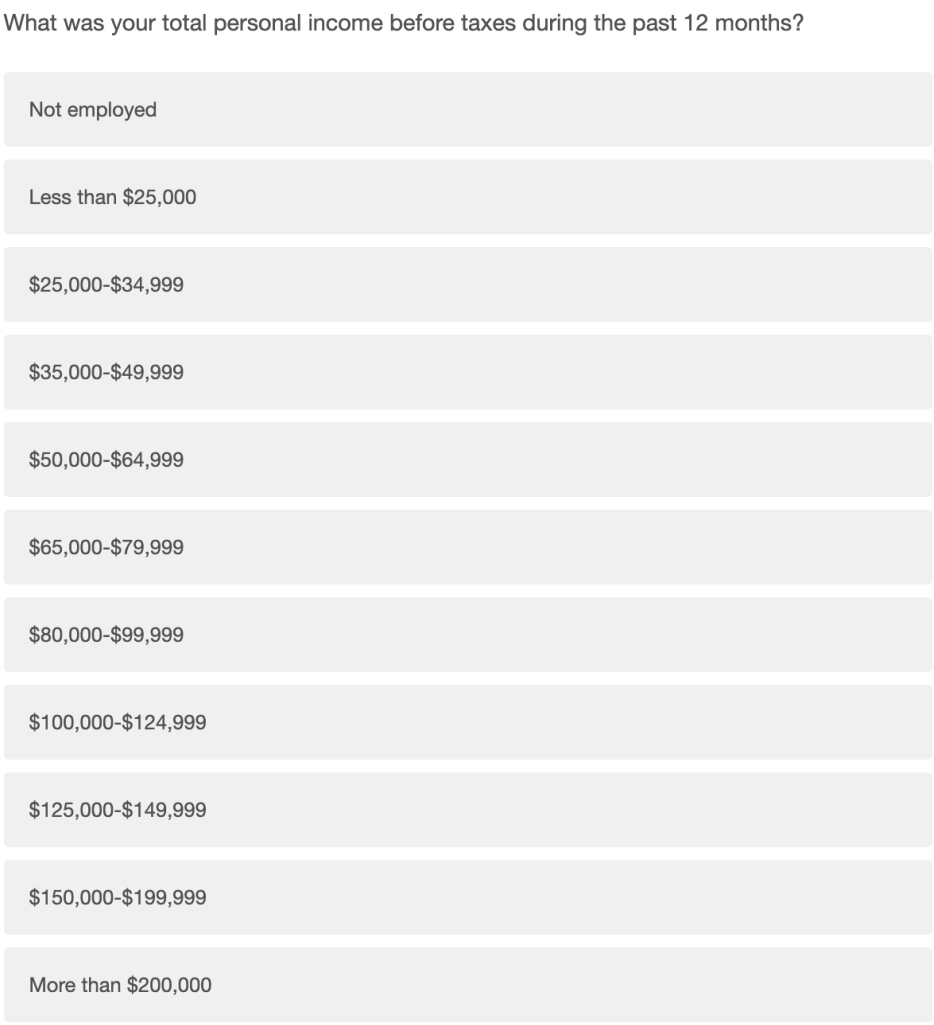
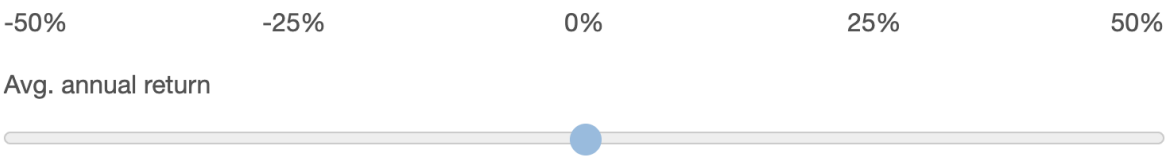


Figure A1: Survey Instrument (Continued)

In the next few questions, we'll ask you to share your **beliefs about future economic outcomes**. There is no right or wrong answer and your response will in no way affect your compensation.

Imagine as a long term investment you invest in the **S&P 500 stock market index fund** (a fund that follows the performance of 500 of the largest publicly traded companies in the U.S). If you held that fund for the next **20 years**, what **average annual return** would you expect?



Imagine as a long term investment you invest in a **U.S. housing market index fund** (a fund that simply follows the entire U.S. housing market). If you held that fund for the next **20 years**, what **average annual return** would you expect?

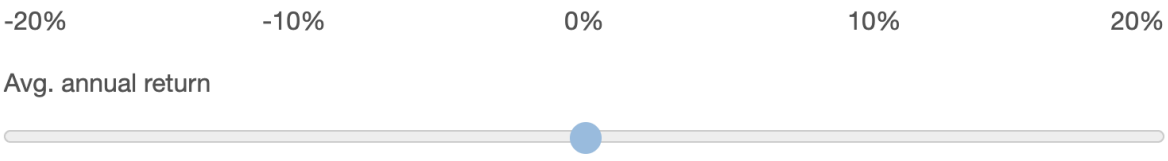


Figure A1: Survey Instrument (Continued)



Figure A1: Survey Instrument (Continued)

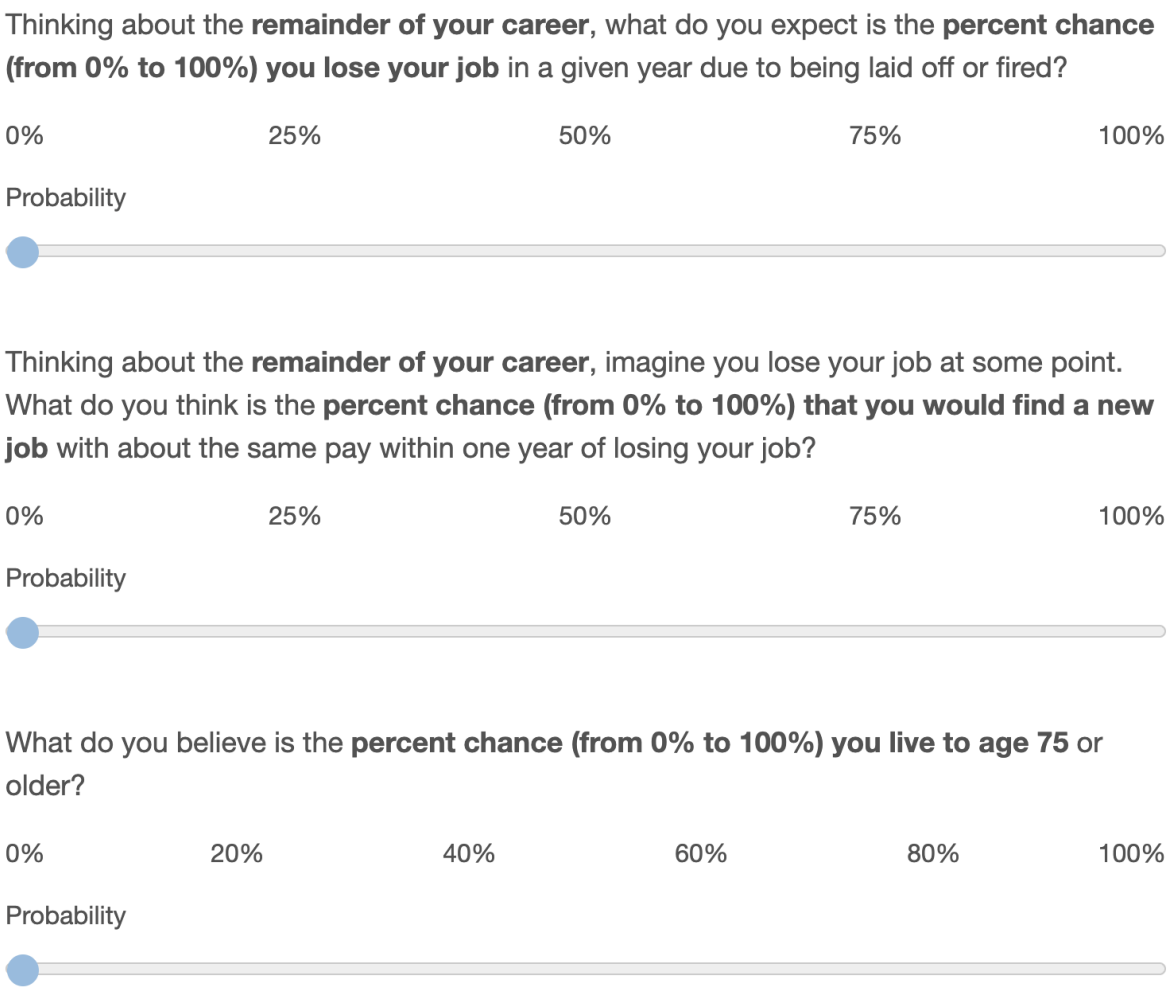


Figure A1: Survey Instrument (Continued)

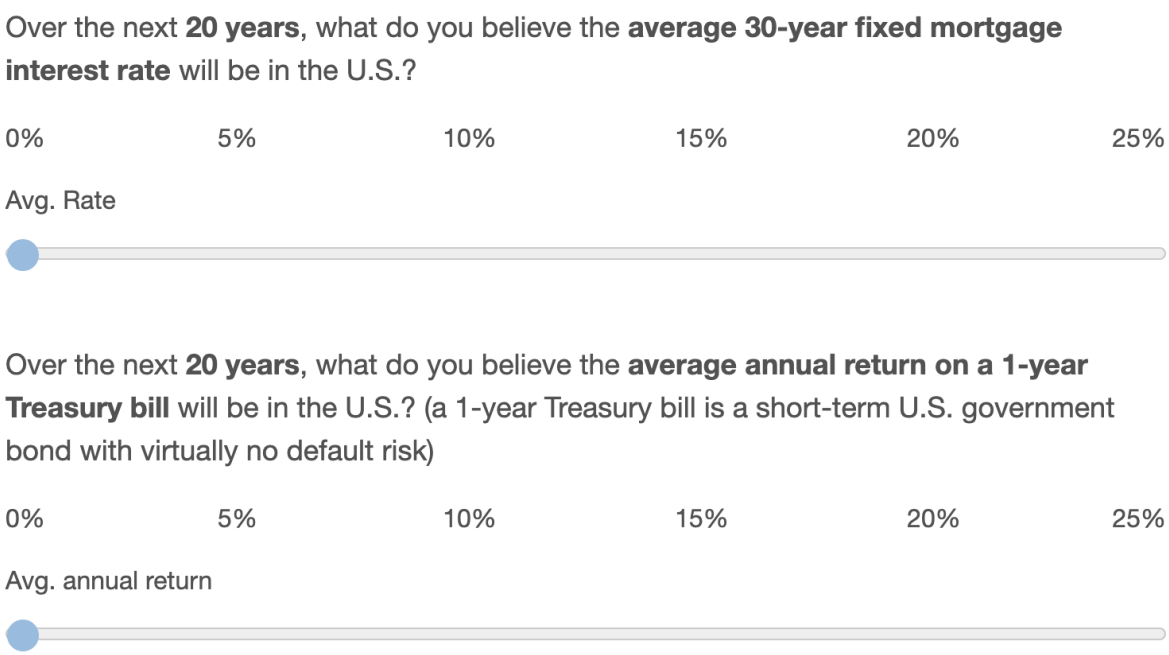


Figure A1: Survey Instrument (Continued)

If the **S&P 500 stock market index were to fall by 10%** over the next year, what do you think would happen to **average U.S. home prices** over the same year?

- home prices fall by more than 10%
- home prices fall by between 5% and 10%
- home prices fall by less than 5%
- Stay about the same
- home prices rise by less than 5%
- home prices rise by between 5% and 10%
- home prices rise by more than 10%
- Unsure

Figure A1: Survey Instrument (Continued)

Think about your own total wages and salary. If the **S&P 500 stock market index were to fall by 10%** over the next year, what do you think would happen to **your wages over the same year** (including any changes from switching jobs or experiencing unemployment)?

- your wages fall by more than 10%
- your wages fall by between 5% and 10%
- your wages fall by less than 5%
- Stay about the same
- your wages rise by less than 5%
- your wages rise by between 5% and 10%
- your wages rise by more than 10%
- Unsure

Figure A1: Survey Instrument (Continued)

Imagine the **S&P 500 stock market index were to fall by 10%** over the next year. **Compared to where prices would have been otherwise**, what do you think **the price of goods and services** after one year would be:

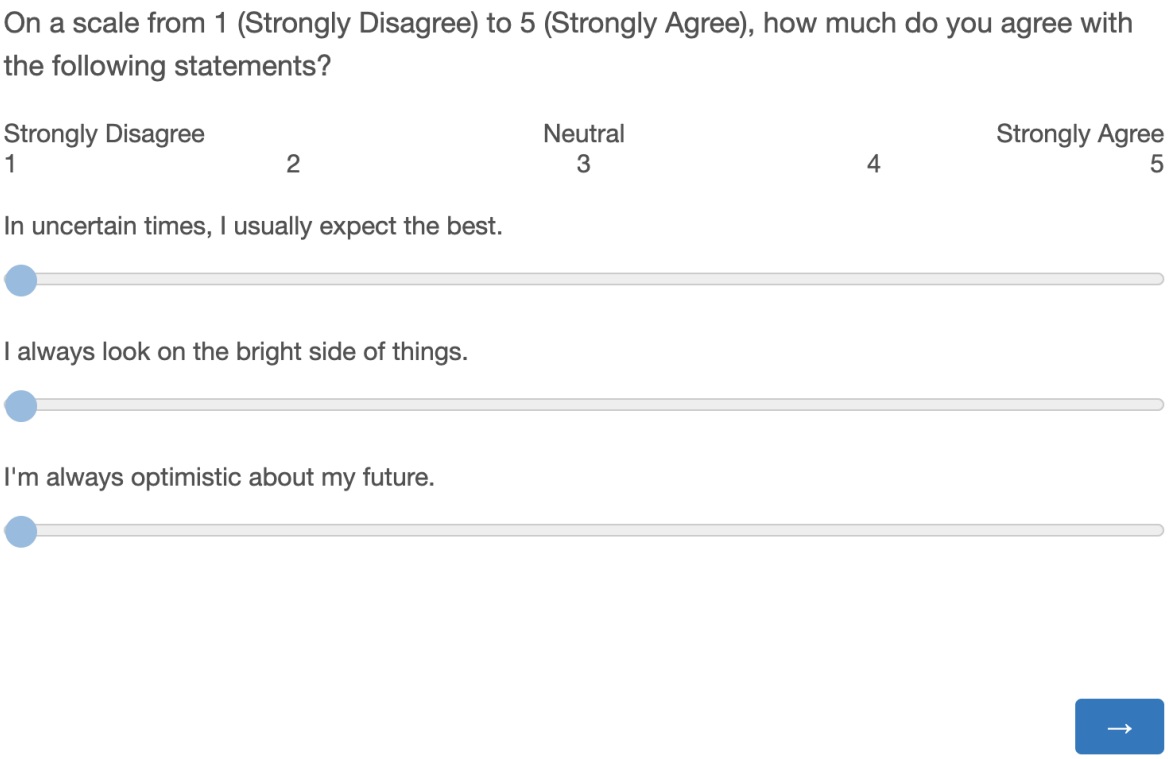
- More than 10% lower
- Between 5% and 10% lower
- Less than 5% lower
- Roughly the same
- Less than 5% higher
- Between 5% and 10% higher
- More than 10% higher
- Unsure

Figure A1: Survey Instrument (Continued)

If the **S&P 500 were to fall by 10%**, what do you think would happen to the **1-year Treasury bill rate**? (A 1-year Treasury bill is a short-term U.S. government bond with virtually no default risk).

- Fall by more than 2 percentage points
- Fall by less than 2 percentage points
- Remain roughly the same
- Rise by less than 2 percentage points
- Rise by more than 2 percentage points
- Unsure

Figure A1: Survey Instrument (Continued)



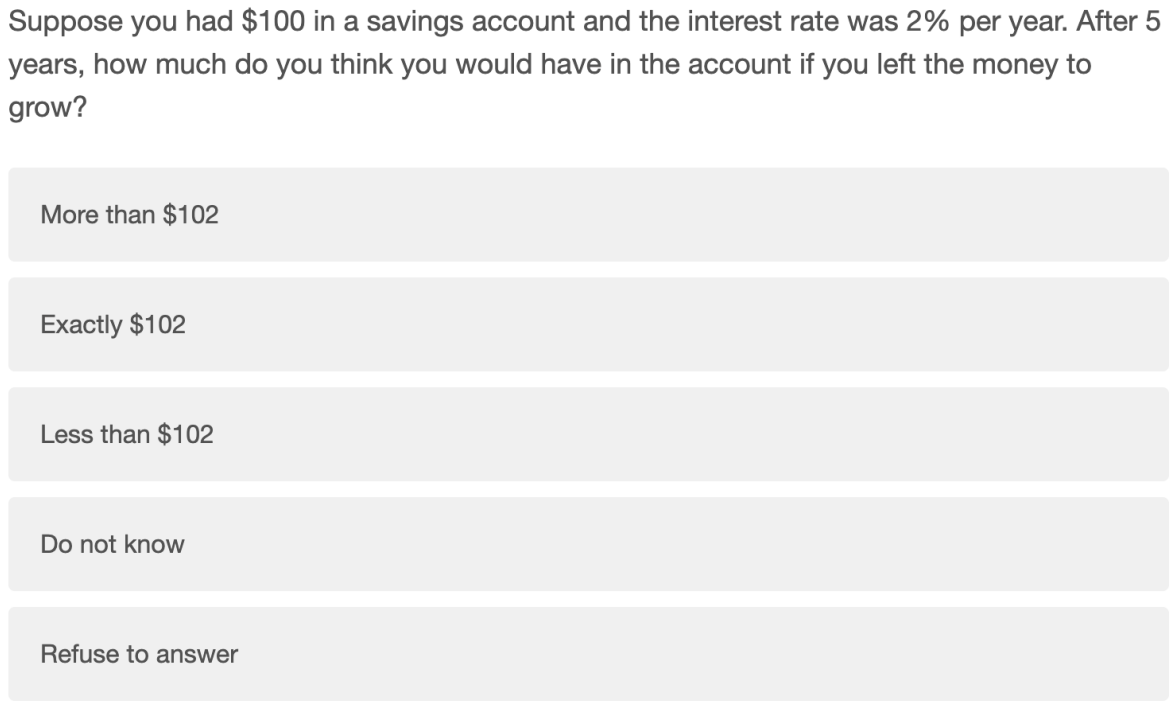


Figure A1: Survey Instrument (Continued)

Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After 1 year, how much would you be able to buy with the money in this account?

- More than today
- Exactly the same
- Less than today
- Do not know
- Refuse to answer

Figure A1: Survey Instrument (Continued)

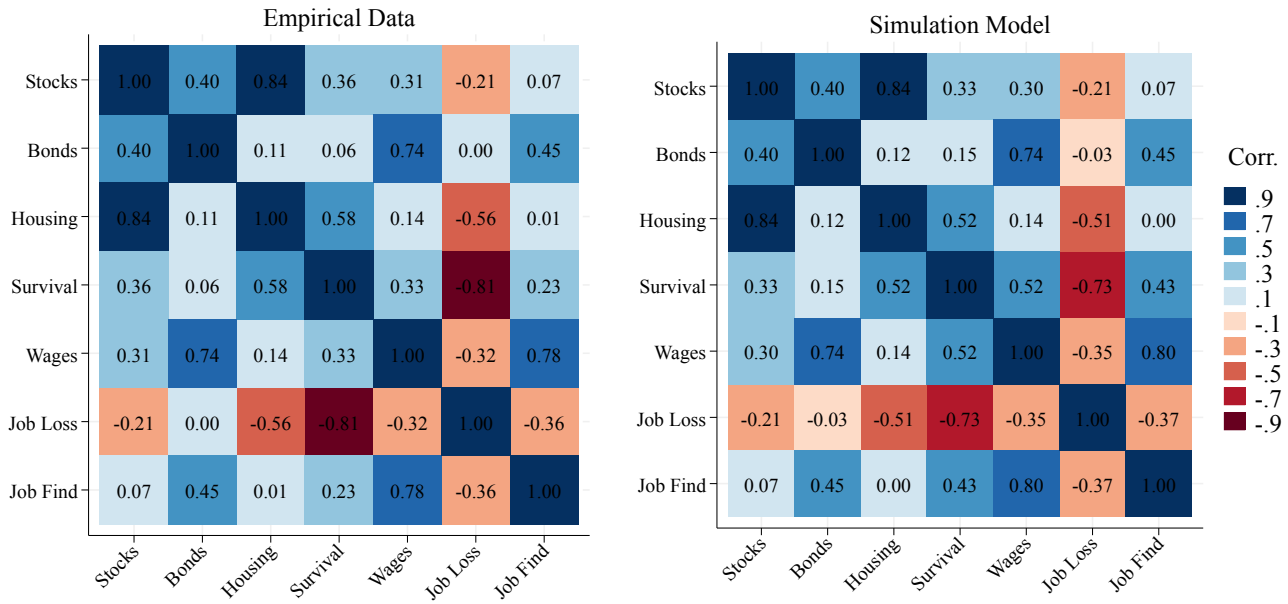
Please tell me whether this statement is true or false. “Buying a single company’s stock usually provides a safer return than a stock mutual fund.”

- True
- False
- Do not know
- Refuse to answer

If interest rates rise, what will typically happen to bond prices?

- They will rise
- They will fall
- They will stay the same
- There is no relationship between bond prices and the interest rate
- Do not know
- Refuse to answer

Figure A2: Empirical vs. Simulated Correlations



Notes: This figure compares the cross-sectional correlation of the seven key economic beliefs used in the model between the survey data (left) and the structural model simulation (right). Empirical moments are calculated from survey respondents who participated in both survey rounds, using the second round for cross-sectional moments. Moments are aggregated by belief group. Beliefs for stocks, bonds, housing, and wages are expressed as real expected growth rates (net of subjective inflation). Job loss, job find, and survival expectations are expressed as subjective probabilities.

B Appendix B: Construction of Common Factors

The purpose of our factor analysis is to identify low-dimensional latent belief constructs that summarize the correlation structure across respondents' expectations. Starting from ten standardized expectation measures, we extract a small number of common factors that capture shared variation in beliefs about asset prices, liabilities, and human capital. Throughout this exercise, we allow the data to determine both the magnitude and sign of factor loadings without imposing any prior interpretation on individual expectations. In particular, the estimation procedure is not informed that high expected stock returns may reflect optimism while high expected inflation may reflect pessimism; instead, the factor structure is inferred entirely from the empirical correlation matrix.

B.1 Common factor model

We rely on a common factor model as described in Fabrigar and Wegener (2012). Let X denote the $N \times p$ matrix of standardized expectations, where p is the number of expectations and N is the number of observations in our data. We standardize all expectations prior to estimation so that factor extraction depends on correlations among beliefs rather than differences in scale or variance.

The common factor model characterizes the correlation structure of observed expectations rather than individual observations. Specifically, the model implies that the variance–covariance matrix of standardized expectations satisfies

$$\mathrm{Var}(X) = \Lambda\Lambda' + D_\psi,$$

where Λ is a $p \times m$ matrix of factor loadings and D_ψ is a diagonal matrix collecting expectation-specific (unique) variances. The (i,k)-th element of Λ, denoted by λ_{ik}, represents the loading of expectation i on latent factor k and measures the strength and direction of their association.

Under this representation, the covariance between two distinct expectations i and j is given by

$$\mathrm{Cov}(X_i, X_j) = \sum_{k=1}^{m} \lambda_{ik}\lambda_{jk}, \qquad \text{for } i \neq j.$$

Thus, expectations are correlated only to the extent that they share exposure to the same latent factors. The variance of expectation i decomposes as

$$\mathrm{Var}(X_i) = \sum_{k=1}^{m} \lambda_{ik}^2 + \psi_i,$$

where ψ_i denotes the unique variance of expectation i. The quantity

$$h_i^2 = \sum_{k=1}^{m} \lambda_{ik}^2$$

is referred to as the *communality* and measures the fraction of variation in expectation i explained by the common factors.

B.2 Estimation and rotation

We estimate the model using principal factor analysis, which extracts factors from the shared (common) variance of the observed variables rather than total variance. The procedure iteratively estimates communalities, replaces the diagonal of the correlation matrix with these values, and extracts the leading eigenvectors until convergence. We retain $m = 3$ factors, which together capture a substantial fraction of the common correlation structure and yield interpretable latent dimensions. To aid interpretation, we apply a promax rotation, an oblique rotation method that allows the extracted factors to be correlated. Allowing for correlated factors is appropriate in our setting, as beliefs about asset returns, interest rates, and labor market outcomes need not be statistically independent. Rotation does not change the model's fit but selects a representation of the factor space that yields a simpler loading structure, with expectations loading strongly on a small number of factors.

B.3 Factor loadings

Figure 2 plots the rotated factor loadings for each expectation using a lollipop-style visualization. For each factor, expectations are ordered by their signed loading values. Larger absolute loadings indicate a stronger association between an expectation and the corresponding latent belief dimension, while the sign reflects the direction of the relationship. Across factors, the loadings reveal distinct belief bundles. One factor is primarily associated with expectations about asset prices and returns, another loads heavily on interest rates and inflation, and a third captures expectations related to labor market outcomes and longevity. These patterns motivate our interpretation of the factors as latent belief constructs governing different domains of economic expectations.