STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
5350 Master's thesis in economics
Spring 2026

Real Term Premia in Life-Cycle Portfolio Choice

Hugo Ryberg (42597)

Abstract: This paper explores how assumptions about real term premia and the resulting time-varying expected excess returns in bonds alter optimal life-cycle portfolio decisions. I construct 10-year real-yield specifications that progressively scale the real term premium from a zero-premium expectations-hypothesis lower bound to a Fisher-adjusted upper bound. These yields are derived by deflating nominal rates with an AR(1) inflation forecast and assuming a zero inflation-risk premium. For each yield specification, I estimate an annual VAR(1) over a 1920–2011 U.S. sample for equity and 10-year real bond returns, incorporating yield-curve and valuation-based predictors. I then embed these estimated return dynamics into a life-cycle portfolio-choice model with risky labor income, Social Security, mortality risk, and a bequest motive. This setup provides a structured counterfactual: how much does the introduction of real term premia shift life-cycle portfolio decisions? I find that moving from the expectations-hypothesis baseline to the Fisher-adjusted yield raises optimal long-bond demand from 8% to 75% at age 45. Rather than substituting for equities, this larger real bond demand is financed by taking short positions in real bills. Decomposing this increase reveals that the majority of it is driven not by the unconditionally higher bond premium, but by intertemporal hedging against fluctuations in future real bond risk premia. Finally, life-cycle forces actively shape the overall level of this bond exposure, causing mid-life and late-retirement households to hold more 10-year real bonds than an infinite-horizon investor.

Keywords: Real term premia; life-cycle portfolio choice; strategic asset allocation; long-term real bonds; return predictability; intertemporal hedging

JEL: G11; G12; E43; D14

Supervisor:	Zoltan Racz
Date submitted:	May 18th, 2026
Date examined:	May 29th, 2026
Discussants:	Elias Paulsson and Oliver Westin Bergh
Examiner:	Sampreet Singh Goraya

Declaration of AI and Computational Tool Usage

The preparation of this thesis was supported by a combination of AI-assisted writing, coding, and computational tools. OpenAI Prism served as the LaTeX editing environment. Claude Code and ChatGPT Codex were used for coding assistance, debugging, document editing, and workflow support. Refine.ink was used to support proofreading and revision. GPU resources purchased through Lambda.ai were used to run the life-cycle model.

All AI-assisted and computational outputs were reviewed, assessed, and edited prior to incorporation. I remain solely responsible for the thesis text, code, analysis, interpretation, and conclusions.

Hugo Ryberg

Contents

1 Introduction

For long-horizon investors, long real bonds play a dual role: they are theoretically the relevant risk-free asset, yet empirically a volatile source of compensated risk. On the theoretical side, Campbell and Viceira (2001) show that an inflation-indexed bond portfolio matched to future consumption flows serves as the true safe asset for a long-horizon investor, delivering known real cash flows that eliminate the reinvestment risk of rolling over short-term bills. On the empirical side, the expectations hypothesis fails to capture the observed behavior of long yields, as investors demand time-varying compensation for holding term risk. A rich literature, including Fama and Bliss (1987), Campbell and Shiller (1991), and Cochrane and Piazzesi (2005), demonstrates that yield spreads and forward rates systematically forecast excess returns on long bonds. This confirms that long yields contain a dynamic term premium alongside expected future short rates, a feature Pflueger and Viceira (2011) corroborate specifically for inflation-indexed markets. Ultimately, these two distinct features—the guaranteed cash flows and the fluctuating risk premium—are bound together by movements in the yield curve. Mechanically, a rise in yields inflicts an immediate capital loss. When the portfolio's real cash flows are matched to future consumption, higher subsequent returns exactly offset this loss, inducing demand for the bond as a safe haven. However, to the extent that this yield increase is driven by a time-varying term premium, it also signals a higher expected excess return. This higher term premium introduces two additional sources of portfolio demand. First, the higher expected reward improves the bond's risk-adjusted return, generating standard static demand to capture the elevated premium. Second, because the immediate capital loss coincides with a rising conditional risk premium, the bond also serves as a powerful intertemporal hedge against shifts in its own expected excess returns. Balancing the need to hold the bond for its long-term safety against the incentive to capture time-varying risk premia, investors face a portfolio problem that hinges on a central question: how much of the term spread should be interpreted as expected future real short rates, and how much as a compensated real term premium?

To address this question, I evaluate how sensitive life-cycle portfolio allocations between short bills, long bonds, and equities are to assumptions about the real term premium in an environment with time-varying expected returns. The portfolio response is driven by two intertwined forces: a higher average compensation for holding term risk, and the predictability of those returns, which dictates how well long bonds hedge their own future excess returns. However, choosing a single, definitive calibration for these return dynamics is difficult because the empirical literature offers weak guidance on the precise magnitude of the term premium. While it is clear that bond risk premia exist and vary over time, estimates diverge widely depending on whether they are inferred from the short histories of inflation-indexed bond markets or backed out of nominal yields. Furthermore, Bauer et al. (2012) show that these estimates are highly sensitive to models of persistent yield dynamics, since expected future short rates must be projected many years forward. Because reliable estimation of these dynamics requires long time series that inflation-indexed bond markets cannot yet provide, the exact scale of the real term premium remains an empirical challenge.

Because the true premium cannot be precisely measured, this thesis analyzes life-cycle behavior across a range of yields defined by two bounds. Both bounds rely on underlying real rates constructed by subtracting an AR(1) inflation forecast from historical U.S. nominal data. The lower bound is the expectations-hypothesis yield $y^{EH}_{10,t}$, a standard assumption widely used in the portfolio-choice literature when constructing bond returns. This baseline imposes the restriction that the real yield on a 10-year bond exactly equals the average expected return of rolling over one-year real bills for 10 years. Under this assumption, there is zero expected excess return to compensate investors for holding term risk. The upper bound is a Fisher-adjusted 10-year yield $y^{FA}_{10,t}$, which serves as a natural ceiling because it attributes the entire sample mean of the yield spread to term premia. By assuming a zero inflation risk premium, it forces any term premium in nominal bonds to be treated as a real term premium. The gap between this Fisher-adjusted yield and the expectations-hypothesis baseline constitutes the full implied real term premium. These different bounds can be formulated as nested specifications dependent on a scaling parameter λ, given by

$$y_{10,t}(\lambda) = y^{EH}_{10,t} + \lambda \underbrace{\left(y^{FA}_{10,t} - y^{EH}_{10,t}\right)}_{\text{Real Term Premium}_t}.$$

By setting $\lambda = 0$, the model strictly imposes the expectations-hypothesis baseline, implying long bonds earn no expected excess return. Setting $\lambda = 1$ assigns the modeled long bond the Fisher-adjusted yield, attributing the full difference between these bounds to the real term premium. I also evaluate $\lambda = 0.5$ to explore an intermediate state. For each λ setting, I estimate a vector autoregression (VAR) that captures how asset returns and economic predictors evolve over time. I then embed each of these λ-dependent financial environments into a life-cycle model where households allocate wealth across equities, long bonds, and short bills. The model incorporates uninsurable background risks and life-cycle features such as risky labor income, progressive Social Security, earnings-dependent mortality, and a bequest motive. Crucially, this setup acts as a structured counterfactual. Because the household's lifetime structure remains constant, shifts in optimal portfolio choice across these environments capture the full effect of the term premium—not only its average size, but the distinct time-series dynamics and return predictability it imparts on the financial environment.

The central finding is that moving from the expectations hypothesis baseline to the Fisher-adjusted real yields substantially shifts the life-cycle portfolio toward long bonds, with the midlife long-bond allocation rising from roughly 8 percent to 75 percent at age 45. I demonstrate that this is not a static response to higher average returns but rather a shift driven by intertemporal hedging. Under the Expectations Hypothesis, expected future bond excess returns are constant. Under the Fisher-adjusted 10-year yield, short-term capital losses on long bonds predictably coincide with upward revisions to future bond risk premia. Holding the long bond therefore allows households to hedge against variation in future expected returns. Furthermore, to capture both this bond premium and the equity premium, households take on substantial leverage by shorting short-term bills.

To substantiate this mechanism, I decompose the 66-percentage-point increase in midlife bond demand. A restricted baseline model with no return predictability shows that just 24 percentage points stem from standard IID demand for the higher average bond return. The remaining 42 percentage points therefore represent an increase in intertemporal hedging demand. To identify which specific variables generate the hedging demand for bonds, I isolate the contribution of each return predictor. Including only the equity valuation predictor in the forecasting model has virtually no effect on the λ-driven increase in bond demand. Instead, the full 42-percentage-point hedging contribution emerges only when the term spread and the short rate are used to forecast returns. Because these yield-curve variables predict future bond returns, this confirms the shift into long bonds is a direct response to the predictability of the real term premium.

It is important to caveat these quantitative findings and explicitly motivate the λ sensitivity analysis. The empirical construction of the 10-year real yield in this paper relies on a standard Fisher adjustment—subtracting a time-series inflation forecast from observed nominal yields. Because this method does not separately isolate an inflation risk premium, the estimated real term premium loaded into the $\lambda = 1$ specification may combine the true real term premium with compensation for inflation risk. To the extent that nominal bonds command a premium strictly for bearing inflation risk, the $\lambda = 1$ specification may overstate the true expected excess return and hedging value of a real long bond. Sweeping the life-cycle model across $\lambda \in \{0, 0.5, 1\}$ directly addresses this empirical ambiguity, mapping specifications from a conservative lower bound containing no real term premium ($\lambda = 0$) up to an upper bound that assigns the entire empirical-minus-EH yield gap to the real bond's dynamics ($\lambda = 1$).

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 formulates the household's dynamic life-cycle optimization problem, including the asset menu, labor income, and bequest motives. Section 4 details the empirical data, the construction of the real yield series, and the estimation of the return-predictability system under varying term-premium assumptions. Section 5 presents the main quantitative results for life-cycle asset allocation across these different yield specifications. Section 6 details the underlying economic mechanisms, compares the findings to prior studies using alternative methods, and explores empirical limitations. Finally, Section 7 concludes.

2 Literature Review

This thesis bridges three literatures. Section 2.1 reviews the asset pricing evidence on term premia to motivate the central empirical exercise: evaluating how the shift from Expectations Hypothesis to Fisher-adjusted long real yields affects portfolio demand. Section 2.2 surveys the dynamic portfolio choice literature, which establishes the intertemporal hedging mechanisms that operate in the model. Due to the computational constraints of adding yield state variables, this literature is typically cast in an infinite-horizon setting. Finally, Section 2.3 reviews the life-cycle portfolio choice literature to motivate the model's calibration and clarify the paper's contribution relative to its closest precedents.

2.1 Term Premia and Interest-Rate Risk

Under the expectations hypothesis, long yields are the average of expected future short rates. Equivalently, excess returns on long bonds should not be forecastable from yield-curve variables. The empirical evidence rejects this prediction. Fama and Bliss (1987) regress one-year excess bond returns on the forward-spot spread, defined as the gap between the n-year forward rate and the current one-year yield. They obtain slope coefficients close to one with an R^2 near 18%. The same forward-spot spread also forecasts changes in the short rate two to four years ahead, with the R^2 rising to 48% at the four-year horizon. This evidence points to both predictable bond excess returns and persistent movements in interest rates. Campbell and Shiller (1991) approach the same hypothesis from the yield-curve side. Regressing future yield changes on current yield spreads, they find slope coefficients well below one and often substantially negative, rejecting the expectations hypothesis. In return space, this implies that high yield spreads forecast positive excess returns on long bonds rather than an offsetting rise in future short rates.

Cochrane and Piazzesi (2005) generalize the predictor. Rather than mapping each bond to its own forward spread, they regress one-year excess returns of all maturities on a joint vector of forward rates. They document that a single tent-shaped linear combination forecasts excess returns on every maturity, yielding an R^2 up to 44%—more than double the Fama-Bliss benchmark. The single-factor structure is robust across maturities, countercyclical, and contains an important component orthogonal to the standard level, slope, and curvature factors that summarize the yield curve. Together, these contributions establish that bond term premia are nonzero, time-varying, and forecastable from yield-curve variables. This motivates including the term spread and short rate as state variables to forecast long-bond excess returns in the model. However, all three studies examine nominal returns. Isolating the real term premium introduces distinct challenges, discussed next.

The object of interest for this study is the real term premium on long real bonds, the part of a long real yield that exceeds the expected average of future real short rates. This is the compensation for bearing long-duration real-rate risk rather than rolling over short real bills. A nominal long bond also exposes a real-consumption investor to inflation risk, since unexpected inflation changes the real value of its fixed nominal payments. The nominal term premium therefore combines compensation for real-rate risk and inflation risk. The real term premium isolates the first component. In empirical term-structure models, however, this object is not observed directly.

Neither nominal yields nor inflation-indexed yields provide a clean measure of the real term premium. Ang et al. (2008) estimate a joint term-structure model of nominal yields, real rates, expected inflation, and inflation risk premia. They demonstrate that nominal yields contain substantial inflation compensation. Thus, while nominal yield predictability confirms the existence of time-varying bond risk premia, it does not isolate the real term premium relevant for a long-horizon investor. While inflation-indexed bonds offer more direct real-yield data, they do not eliminate the measurement problem. Campbell et al. (2009) emphasize that Treasury Inflation-Protected Securities (TIPS) yields reflect expected future real short rates, real risk premia, and liquidity or market-segmentation premia. Furthermore, the inflation-indexed market is too young. The available time series are too short to reliably estimate the forecasting vector autoregression required for the model in this paper. Pflueger and Viceira (2011) test the expectations hypothesis directly in inflation-indexed bond markets. They reject the real expectations hypothesis and document predictable excess returns on real bonds, yet caution that real bond premia remain difficult to separate from liquidity and other market-specific components.

A further caveat is that term-premium estimates are sensitive to finite-sample bias. Stambaugh (1999) shows that predictive regressions with persistent regressors can produce biased and imprecise estimates in samples of realistic length. This problem is closely related to term-structure estimation because expected future short rates depend heavily on estimated persistence. Bauer et al. (2012) make this point directly for dynamic term-structure models. They show that standard estimates of highly persistent yield factors are biased in small samples, and that correcting this bias can materially change estimated expectations components and term premia. To gauge this uncertainty, Section 4 employs a block bootstrap procedure. However, the life-cycle portfolio optimization in Section 5 abstracts from parameter uncertainty and treats the VAR point estimates as known.

These measurement problems motivate the empirical design of this thesis. Whether inferred from nominal yields or extracted from inflation-indexed bonds, the real term premium remains a residual object dependent on assumptions about expected short rates, inflation risk, and liquidity. Because this component is imperfectly identified, the analysis does not rely on a single point estimate. Instead, it studies the portfolio implications through a sensitivity exercise. Specifically, the model evaluates how asset demand changes as long real yields shift from the Expectations Hypothesis to a Fisher-adjusted specification.

2.2 Dynamic Portfolio Choice and Return Predictability

Return predictability renders portfolio choice dynamic. When expected excess returns and term premia vary over time, static mean-variance analysis (Markowitz, 1952) is insufficient. The one-period framework misses two features. First, bond risk depends on the investment horizon: yield changes generate short-run capital gains but also alter reinvestment opportunities and expected future returns. Second, investors finance lifetime consumption rather than terminal wealth. Merton (1971) formalize joint consumption and portfolio choice, showing that investors hold assets both for their short-run return and to intertemporally hedge shifts in investment opportunities.

Implementing this dynamic framework requires specifying empirical variables that forecast returns and solving the resulting optimization problem. Early implementations simplify the problem by assuming the investor maximizes only terminal wealth. Kim and Omberg (1996) derive analytical solutions for a single risky asset, while Brennan et al. (1997) use numerical methods to allocate across stocks, bonds, and cash. Both show that time-varying expected returns generate intertemporal hedging demands, but their focus on terminal wealth assumes away intermediate consumption. Barberis (2000) documents horizon effects on portfolio demand driven by dividend-yield predictability.

Subsequent research introduces intermediate consumption, demonstrating that predictability makes portfolio choice heavily dependent on the investment horizon. Wachter (2002) solves a joint consumption and portfolio problem, showing that mean-reverting stock returns cause optimal equity allocations to increase with horizon. To solve dynamic portfolio problems without full numerical modeling, Campbell and Viceira (1999) develop analytical approximations. Applying these methods, Campbell and Viceira (2001) show that long-term real bonds hedge real interest-rate risk: holding long bonds locks in yield and stabilizes future consumption, making the inflation-indexed perpetuity the benchmark risk-free asset. Extending these horizon effects, Brennan and Xia (2002) and Sangvinatsos and Wachter (2005) introduce richer term-structure dynamics, showing that bond allocations depend critically on the persistence of real rates and time-varying term premia. The common implication across this literature is that return predictability makes optimal asset allocation horizon-dependent. However, these models treat the investment horizon as an exogenous parameter and abstract from labor income.

Campbell et al. (2003) provide the methodological precedent for the return-state modeling used in this thesis. They estimate a first-order vector autoregression (VAR) governing asset returns and predictive state variables. Their baseline specification allocates wealth across bills, nominal bonds, and equity. Solving the portfolio problem for an infinitely lived Epstein-Zin investor, they show that predictability generates substantial intertemporal hedging demand. This demand arises because unexpected asset returns covary negatively with innovations to the state variables that positively forecast future returns. For example, the strong historical covariance between stock returns and dividend-yield innovations raises optimal long-run equity allocations. On the fixed-income side, optimal allocations to nominal bonds vary drastically across historical samples depending on inflation risk and the prevailing stock-bond correlation. To evaluate fixed income absent inflation risk, they construct a hypothetical real consol. Using the VAR-implied path of expected future real short rates, they build this real bond under the expectations hypothesis. Introducing this synthetic asset demonstrates that conservative long-horizon investors allocate heavily to indexed bonds to hedge real interest-rate risk.

These models highlight a gap in the literature. While predictability makes optimal portfolios horizon-dependent, the relevant horizon for households is not an exogenous parameter; it declines naturally with age and interacts with background risks. This thesis adopts the empirical return-state VAR of Campbell et al. (2003) but departs in two ways. First, rather than applying infinite-horizon approximations, it embeds the VAR inside a structural life-cycle model. Second, Campbell et al. (2003) construct their real bond under the restrictive expectations hypothesis. This thesis relaxes this assumption by considering alternative specifications that scale the gap between the expectations-hypothesis benchmark and Fisher-adjusted real yields. This approach captures how time-series variation in the term premium affects life-cycle bond demand.

2.3 Life-Cycle Portfolio Choice

Unlike the infinite-horizon setting examined above, life-cycle models introduce incomplete markets, most notably through uninsurable and non-tradable labor income. Building on Merton (1971), this human capital acts as an implicit background asset. Viceira (2001) shows that when labor-income innovations are weakly correlated with aggregate stock returns, human capital behaves as an implicit risk-free bond. Young investors offset this background exposure by tilting their explicit financial portfolios toward equity. Cocco et al. (2005) provide the first fully specified numerical solution to this incomplete-markets problem, incorporating mortality, borrowing constraints, and idiosyncratic income shocks. Subsequent research refines how uninsurable risk evolves over the business cycle. Due to convex marginal utility, investors are highly sensitive to the higher-order moments of these income shocks. Storesletten et al. (2004) argue that the cross-sectional variance of idiosyncratic income is countercyclical. However, using administrative data, Guvenen et al. (2014) show that variance is roughly constant over the cycle, whereas left-skewness is strongly countercyclical. Catherine (2022) incorporates this cyclical skewness into a life-cycle model. He demonstrates that because uninsurable income tail risk coincides with stock market declines, young households optimally reduce their equity allocations.

Beyond its role as a bond-like asset influencing equity demand, human capital acts as a non-tradable, long-duration claim exposed to interest-rate risk (Munk and Sørensen, 2010; Koijen et al., 2010). Households seek to hedge their future consumption against these interest-rate fluctuations. While perfect immunization requires a full spectrum of zero-coupon bond maturities, the restricted asset menu in standard life-cycle models forces households to approximate a safe position through duration matching. Catherine et al. (2024) show how this duration matching interacts with implicit background assets. Because human capital and Social Security cannot be traded, a household's inherent interest-rate exposure changes mechanically over the life cycle. The background exposure to these long-duration claims is U-shaped: heavily driven by human capital when young, dropping in middle age as human capital depreciates, and rising again in retirement through Social Security. To keep the duration of total wealth matched to the consumption liability, demand for the traded long bond must offset this U-shape. Optimal bond allocations are therefore hump-shaped: low for young investors, peaking in middle age as the background duration falls, and declining as Social

Security assumes the long-duration role.

The mechanisms described above assume constant expected excess returns. Empirical return predictability transforms these baseline age profiles, altering the overall level of asset demand through intertemporal hedging and making optimal allocations depend on financial state variables in addition to age and wealth. The literature explores these dynamics primarily on the equity side. Michaelides and Zhang (2017) and Gomes et al. (2022) show that time-varying equity premia render pure age-based rules, such as standard target-date funds, suboptimal because optimal allocations must condition on predictive state variables. Michaelides and Zhang (2022) confirm that while imperfect predictors dampen this state dependence, they do not eliminate it. Crucially, intertemporal hedging interacts directly with life-cycle wealth dynamics. Duarte et al. (2021) demonstrate that optimal allocations depend jointly on age, wealth, dividend-price ratios, and business-cycle conditions. Furthermore, under incomplete markets, some financial predictors also forecast labor-income dynamics. Lynch and Tan (2011) document that the dividend yield predicts both stock returns and the distribution of labor-income growth. Within the life cycle, dynamic portfolio choice therefore reflects both intertemporal hedging against changing investment opportunities and responses to time-varying uninsurable income risk.

The focal economic mechanism evaluated in this thesis is the role of time-varying term premia for a life-cycle investor within incomplete markets. Previous research typically imposes structural restrictions to make this high-dimensional problem tractable. Sangvinatsos and Wachter (2005) show that time-varying bond risk premia generate substantial intertemporal hedging demands, but abstract from life-cycle frictions such as uninsurable labor income and retirement. Koijen et al. (2010) provide the closest life-cycle analogue. However, they force all wealth into an inflation-linked annuity at retirement, precluding explicit post-retirement portfolio choice. They also restrict return predictability strictly to the bond market and impose borrowing constraints that limit the use of bond-timing strategies early in life. This thesis abandons these restrictions to evaluate dynamic portfolios over the entire life span. It explicitly models portfolio choice through retirement and accommodates joint predictability, utilizing a system where state variables simultaneously forecast equity and bond returns. To reflect the portfolio capabilities of delegated fund managers and target-date funds with derivative overlays, it relaxes retail borrowing limits and permits balance-sheet expansion through short positions in bills. Finally, rather than moving across states within a fixed term-structure model, this thesis introduces an explicit parameter, λ. By scaling the fraction of long-yield variation attributed to term premia while holding uninsurable background risk constant, this approach cleanly separates the age-dependent duration profile from the dynamic compensation for term-premium risk.

3 Model

This section details the household dynamic programming problem and the returns and investment opportunities. The model's life-cycle features closely follow Catherine et al. (2024), while the return dynamics build on the VAR(1) framework of Campbell et al. (2003). However, the implemented specification departs from both papers to reduce the dimension of the state space while retaining the features needed to study the effect of term premia for household portfolio choice. To economize on state variables, the model does not explicitly track Average Indexed Monthly Earnings (AIME) used in pension calculations, unlike Catherine et al. (2024). Instead, the persistent income state serves as a proxy for lifetime earnings. Furthermore, compared to Campbell et al. (2003), the return specification restricts realized returns from acting as forecasting variables and simplifies the asset menu to include only real fixed-income assets and equity. The life-cycle parameters are reported in Appendix A, Table A.1.

3.1 Utility

I model a household with CRRA preferences, age- and income-dependent mortality, and an annuity-equivalent bequest motive. The agent works from age 22 through age 66, retires at age 67, and faces a terminal age of 99. At the beginning of each period, the household observes cash-on-hand X_t, the persistent income state z_t, and the financial state $\mathbf{s}_t$. The dynamic program is written conditional on a fixed return system; for such a system, $\mathbf{s}_t = (\text{cape}_t, spr_t, y_{1,t})'$, where cape_t is the CAPE-based log earnings yield, $-\log(\text{CAPE}_t)$, $y_{1,t}$ is the one-year real yield, and spr_t is the real 10-year term spread associated with that return system. Section 4 constructs the long real yield, and hence the term spread, separately for each value of λ; once a particular return system is fixed, I suppress the λ argument. Cash-on-hand includes income already received at the start of the period. The household cannot borrow against future income. It chooses consumption C_t subject to $0 < C_t \leq X_t$ and portfolio shares α_t over nonnegative invested savings $A_t = X_t - C_t$. The level variables X_t, C_t, A_t, and Y_t are measured in AWI units; lower-case returns and yields are in logs. Its problem is

$$\begin{aligned} V_t(X_t, z_t, \mathbf{s}_t) = \max_{0<C_t\leq X_t, \alpha_t} \quad & u(C_t) + \beta\, \mathbb{E}_t\big[\psi(t, z_t)\, V_{t+1}(X_{t+1}, z_{t+1}, \mathbf{s}_{t+1}) \\ & + \big(1 - \psi(t, z_t)\big)\, b\big(A_t R^p_{t+1}, \mathcal{A}(\mathbf{s}_t)\big)\big], \end{aligned} \tag{1}$$

where $u(C) = C^{1-\gamma}/(1-\gamma)$, R^p_{t+1} is the gross portfolio return on invested savings, and $\psi(t, z)$ is the probability of surviving from age t to t+1. Conditional on survival, next-period cash-on-hand is

$$X_{t+1} = A_t R^p_{t+1} + Y_{t+1}. \tag{2}$$

Conditional on death between periods t and t+1, income is not received. The household derives bequest utility from the invested estate $A_t R^p_{t+1}$ left to heirs. Annuity conversion terms are fixed at the decision date, so the estate is translated into annuity-equivalent consumption using current yields. Following Catherine et al. (2024), I specify the bequest motive as a preference over heirs' annuity consumption. An estate W is converted into a constant annual consumption flow for heirs over $\bar{b}$ years, and the household values that annuity flow with CRRA utility:

$$b(W, \mathcal{A}) = \bar{b} \cdot \frac{(W/\mathcal{A})^{1-\gamma}}{1-\gamma}, \qquad \mathcal{A}(\mathbf{s}_t) = \sum_{k=1}^{\bar{b}} \exp(-k y_{k,t}), \qquad y_{k,t} = y_{1,t} + \frac{k-1}{\bar{b}-1} spr_t, \tag{3}$$

where $W/\mathcal{A}(\mathbf{s}_t)$ is the annual annuity flow implicit in the estate. The parameter $\bar{b}$ is the bequest horizon. The exponential discount factors convert the interpolated log yields into annuity prices using the current yield-state variables.

The one-year death probability for an agent of age t in persistent-income state z is

$$m(t, z) = \min\big\{\chi(z)\, m_{\text{base}}(t),\, 1\big\}, \qquad \psi(t, z) = 1 - m(t, z). \tag{4}$$

The baseline $m_{\text{base}}(t)$ is the gender-averaged SSA 2017 period life table. The scalar $\chi(z)$ is an age-independent multiplier on baseline mortality, below one for high earners and above one for low earners. I calibrate it pointwise on the z-grid using the income-gradient in life expectancy from Chetty et al. (2016). Each grid point z_i is converted into a Gaussian percentile rank using the standard normal cumulative distribution function, $\Phi(z_i/\sigma_z)$, where σ_z is the stationary standard deviation of z.[1] I map this rank to the corresponding income percentile in Chetty et al. (2016), with endpoint clipping, and choose $\chi(z_i)$ so the model matches the associated age-40 life-expectancy target.

3.2 Labor

Working-age labor income combines a deterministic age profile with persistent and transitory components whose shocks are mixture-normal, calibrated to match the skewness and kurtosis of U.S. earnings data documented by Guvenen et al. (2021). Gross labor income is measured in units of the Social Security Average Wage Index, $\bar{L} \equiv 1$, and log income is governed by

$$\log L_t = g_t + z_t + \varepsilon_t. \tag{5}$$

The deterministic age component g_t is a cubic polynomial in age, calibrated as in Catherine et al. (2024) to the average log-earnings profile of U.S. male workers in SSA administrative data. The intercept normalization applies to the full stochastic earnings process, not to the deterministic age profile alone. The persistent component z_t follows an AR(1) with mixture-normal innovations,

$$z_t = \rho\, z_{t-1} + \eta_t, \qquad \eta_t \sim \mathrm{NM}\big(p_z,\, \mu_{\eta 1},\, \sigma^2_{\eta 1},\, \mu_{\eta 2},\, \sigma^2_{\eta 2}\big)\,.^2 \tag{6}$$

The process is highly persistent ($\rho = 0.991$), and the mixture in η combines frequent small positive drifts with rare large negative shocks, producing strong left-skewness. The transitory component ε_t is i.i.d. mixture-normal,

$$\varepsilon_t \sim \mathrm{NM}\big(p_\varepsilon,\, \mu_{\varepsilon 1},\, \sigma^2_{\varepsilon 1},\, \mu_{\varepsilon 2},\, \sigma^2_{\varepsilon 2}\big)\,. \tag{7}$$

The transitory density has strong kurtosis from a rare wide component. The resulting innovations to the labor income process exhibit both left-skewness and high kurtosis.

Working-age gross income is subject to two levies. A Social Security payroll tax applies at rate τ_{SS} to gross income capped at $\bar{y}_{\mathrm{cap}}$,

$$T_{\mathrm{SS}}(L_t) = \tau_{\mathrm{SS}} \cdot \min\{L_t,\, \bar{y}_{\mathrm{cap}}\}. \tag{8}$$

The rate $\tau_{\mathrm{SS}} = 10.6\%$ corresponds to the Old-Age and Survivors Insurance portion of FICA. The Disability Insurance component is excluded, consistent with the absence of disability risk in the model. The remainder is subject to a stylized progressive federal income tax $T_{\mathrm{inc}}(\cdot)$, based on the 2019 TCJA single-filer marginal rates with seven brackets reported in Appendix Table A.2. For modeling tractability, this stylized schedule omits the statutory standard deduction and applies directly to income net of the payroll tax. Disposable income,

$$Y_t = L_t - T_{\mathrm{SS}}(L_t) - T_{\mathrm{inc}}\big(L_t - T_{\mathrm{SS}}(L_t)\big)\,, \tag{9}$$

enters the budget constraint.

[1] Although η is mixture-normal, persistence $\rho = 0.991$ shrinks its standardized skew, leaving the stationary distribution of z approximately Gaussian (skew ≈ -0.15, excess kurtosis $\approx +0.02$).

[2] In both mixtures, $\mu_{\cdot 2}$ is pinned by the zero-mean condition $\mu_{\cdot 2} = -\frac{p}{1-p}\,\mu_{\cdot 1}$, so the persistent and transitory innovations are mean-zero in logs.

At retirement, labor income is replaced by a Social Security pension. The benefit follows the U.S. Primary Insurance Amount (PIA) schedule reported in Appendix Table A.3, applied to an annualized approximation of Average Indexed Monthly Earnings (AIME). Annualized AIME is approximated by the persistent state at retirement scaled by the average age-earnings profile,

$$\text{AIME}(z_R) = \min\{\exp(z_R) \cdot \bar{g},\, \bar{y}_{\text{cap}}\}\,, \tag{10}$$

where $\bar{g} = \overline{\exp(g_t)}$ over working years. The approximation avoids introducing an additional state variable and is justified by the persistence of z. Pension benefits are taxed under the same progressive income tax T_{inc} that applies to labor income, but are exempt from payroll tax.

3.3 Asset Returns and Portfolio Decision

The household can invest in three assets: a one-period real bill, aggregate equity, and a long real bond. The investment opportunity set is summarized by the equity valuation state, the real term spread, and the one-year real yield. All returns are measured in logs. The real bill is conditionally risk-free and pays

$$r^f_{t+1} = y_{1,t}, \tag{11}$$

where $y_{1,t}$ is the one-year real yield observed at the beginning of period t. Let $\mathbf{x}_{t+1} = (x^s_{t+1}, x^b_{t+1})'$ denote excess log returns on equity and the long bond over the bill. The household chooses $\alpha_t = (\alpha^s_t, \alpha^b_t)'$, with residual share $1 - \mathbf{1}'\alpha_t$ invested in the real bill. The model imposes no borrowing constraints or short-selling limits on α_t. This unconstrained setup is intended to reflect the capability of modern target-date funds and institutional pension providers to deploy leverage on behalf of households.

Asset returns and state variables are governed by a Gaussian VAR(1). To make the dynamic program tractable, I impose that lagged realized excess returns do not forecast next-period state variables or excess returns. The VAR innovations may nevertheless be contemporaneously correlated, so shocks to the state variables can be associated with shocks to realized excess returns within the same period. The restricted joint system is

$$\begin{pmatrix} s_{t+1} \\ \mathbf{x}_{t+1} \end{pmatrix} = \begin{pmatrix} \Phi^s_0 \\ \Phi^x_0 \end{pmatrix} + \begin{pmatrix} \Phi_{ss} \\ \Phi_{xs} \end{pmatrix} s_t + \begin{pmatrix} v^s_{t+1} \\ v^x_{t+1} \end{pmatrix}, \qquad \begin{pmatrix} v^s_{t+1} \\ v^x_{t+1} \end{pmatrix} \sim \mathcal{N}\left(0, \begin{pmatrix} \Sigma_{ss} & \Sigma_{sx} \\ \Sigma_{xs} & \Sigma_{xx} \end{pmatrix}\right). \tag{12}$$

Given this return system, the portfolio log return is approximated using the discrete-time approximation to a continuously rebalanced portfolio in Campbell and Viceira (1999),

$$r^p_{t+1} = y_{1,t} + \alpha_t'\mathbf{x}_{t+1} + \frac{1}{2}\alpha_t' \left(\text{diag}(\Sigma_{xx}) - \Sigma_{xx}\alpha_t\right), \tag{13}$$

where the final term is the Jensen correction for return volatility. With this specification, next-period cash-on-hand is

$$X_{t+1} = A_t \exp(r^p_{t+1}) + Y_{t+1}. \tag{14}$$

Since $\exp(r^p_{t+1}) > 0$, return realizations do not generate negative financial wealth under this approximation.[3]

[3]This positivity property differs from exact discrete-time linear aggregation of leveraged asset returns, under which sufficiently adverse return realizations can imply nonpositive financial wealth. Evaluating exact linear returns at the policy points underlying the reported comparative statics indicates that this distinction is immaterial over the reported policy region, although it can matter in remote, highly leveraged tail-state cells.

The term spread spr_t is the key state variable for the paper's sensitivity exercise. It summarizes the long-rate component of the return system and governs the bond-return predictability faced by the household. The calibration section describes how the term-premium content of the long real yield is varied and how this changes the implied spread and long-bond excess-return process.

3.4 Solution Method

The problem admits no closed-form solution. I solve it by backward induction on a discretized state space, using the endogenous grid method of Carroll (2006). At the terminal age $T = 99$, the continuation value is set to zero, so savings carried beyond the terminal period generate only bequest utility:

$$V_T(X_T, z_T, \mathbf{s}_T) = \max_{0<C_T\leq X_T,\alpha_T} \left\{ u(C_T) + \beta\, \mathbb{E}_T\big[b\big((X_T - C_T)R^p_{T+1}, \mathcal{A}(\mathbf{s}_T)\big)\big] \right\}. \tag{15}$$

At each earlier age, state and savings cell, I first solve the two portfolio first-order conditions jointly by a Newton–Raphson iteration and backtracking line search, and then invert the consumption Euler equation to recover the implied consumption. The implied endogenous cash-on-hand is then mapped onto the exogenous cash-on-hand grid by linear interpolation.

The persistent and transitory labor income innovations are two-component normal mixtures, so Rouwenhorst's grid-based transition method for AR(1) processes with Gaussian innovations does not apply. I instead place persistent income on an evenly-spaced grid covering approximately 99% of its unconditional probability mass. To evaluate expectations, I apply a Gaussian quadrature rule tailored to the mixture densities of the persistent and transitory innovations.

The financial state $\boldsymbol{s}_t$ lives on a tensor-product grid built by Cholesky-factoring the stationary VAR covariance $\Sigma_s = LL'$, placing uniformly-spaced points in each standardized direction. I cover approximately 99% of the joint stationary mass. For integration over the VAR innovation $v_t \sim \mathcal{N}(0, \Sigma_v)$, I Cholesky-factor Σ_v and place a tensor-product Gauss–Hermite rule on the orthogonalized innovations, preserving the joint covariance structure while keeping the quadrature one-dimensional on each axis.

Because the choice at any specific grid point relies only on the known next-period value function, the per-cell Newton solves are entirely independent across the discretized state space. This conditional independence makes each age step embarrassingly parallel. I implement the core solution algorithm in JAX, using `vmap` to vectorize the operations across the entire state grid. The routine is then `jit`-compiled for performance and sharded across four NVIDIA H100 GPUs.

4 Estimation

To capture the return-predictability and intertemporal hedging motives central to the life-cycle problem, the model requires an empirical return-state system that governs the investment opportunity set. Estimating the persistent dynamics of interest rates, equity valuations, and term premia necessitates a very long historical sample, rendering modern, short-dated inflation-indexed bond data insufficient. However, constructing real yields over a century of historical nominal data inherently introduces measurement noise and makes the perfect identification of the real term premium impossible. This constraint motivates the paper's core sensitivity approach. This section details that empirical strategy. Section 4.1 describes the underlying 1920–2011 annual dataset used to overcome the sample-length constraint. Section 4.2 outlines the construction of the real interest-rate process and introduces the scaling parameter λ used to introduce real term premia. Section 4.3 presents the estimation of the VAR(1) return system under these varying expectations-hypothesis assumptions, and Section 4.4 interprets the resulting dynamics through the lens of intertemporal hedging demand.

4.1 Data and Variable Construction

The data source used throughout this paper is the updated and revised data associated with Shiller (1989), accessed from Shiller's website. I use the annual figures from Shiller (1989), which provide January-of-year observations of the S&P composite stock price index (`P`), the associated dividend (`D`), the one-year nominal Treasury rate (`R`), the 10-year nominal Treasury yield (`RLONG`), and the consumer price index (`CPI`). I supplement this with Shiller's monthly file, which provides monthly observations of the consumer price index and the Cyclically Adjusted Price-Earnings ratio (`CAPE`). From these raw series I construct the primitive log quantities that enter the analysis. The nominal short and long Treasury yields are expressed in log returns, $y^N_{1,t} = \log(1 + \mathtt{R}_t/100)$ and $y^N_{10,t} = \log(1 + \mathtt{RLONG}_t/100)$. I define the valuation state cape_t as the CAPE-based log earnings yield, $\mathrm{cape}_t = -\log(\mathtt{CAPE}_t)$; hence larger values correspond to cheaper equity valuations. From the monthly consumer price index, I construct the realized December-over-December log inflation series. To align with the January-of-year-t information set used for portfolio decisions, I define $\pi_t = \log(\mathrm{CPI}^{\mathrm{Dec}}_{t-1}/\mathrm{CPI}^{\mathrm{Dec}}_{t-2})$. This relies on standard conventions: because the BLS releases each month's CPI with a roughly three-week lag, the December $t-1$ CPI is the most recent fully published CPI observation available in January of year t. The series π_t therefore represents the realized twelve-month log inflation over calendar year $t-1$ that is observable at the start of year t.

The final sample is annual, sampled at January of each year, and runs from 1920 to 2011 inclusive, yielding $T = 92$ observations. The 1920 start corresponds to the consolidation of active U.S. monetary policy following the post-WWI transition of the Federal Reserve from a Treasury-financing institution to an independent monetary authority, as described by Meltzer (2003) and Friedman and Schwartz (1963). The 2011 endpoint is fixed by data availability: the one-year nominal Treasury rate `R` in the dataset of Shiller (1989) has its last non-missing observation in 2011. Even this conservative sample is not splice-free for the interest rate series. The one-year rate pre-1929 is constructed from commercial paper, since the U.S. Treasury did not auction Treasury bills until December 1929. After 1929 it is taken from Treasury rates. The long-term interest rate pre-1953 is a composite series, and post-1953 it is the 10-year Treasury constant-maturity yield (`GS10`). Two notable episodes fall outside this sample. The pre-1920 period contains the WWI inflation shock, inflation peaked at +18.6 percentage points in 1918 and +13.6 percentage points in 1919, as well as the long deflation of the late nineteenth century and the pre-Federal-Reserve banking panics of 1873, 1893, and 1907. The post-2011 period contains the secular bull market in U.S. equities during which the S&P composite roughly quadrupled, the 2021–2023 inflation spike that followed the COVID-19 pandemic, and the following rapid tightening of monetary policy.

A number of alternative data sources are available for U.S. stock and bond data, but all suffer from short-dated samples. For stock returns, the Center for Research in Security Prices (CRSP) provides a more granular time series from 1926 onwards, and is the standard source in the asset-pricing literature. For interest rates, the Federal Reserve Economic Database (FRED) provides individual yield series at higher frequency, including the three-month and one-year Treasury rates (`TB3MS`, `GS1`) and the 10-year Treasury constant-maturity yield (`GS10`), though most of these series do not extend before 1953. Inflation-indexed Treasury yields (`DFII10`) provide a direct measure of the real long yield but are available only from 1997. The U.K. counterpart, Index-linked Gilts, provides a slightly longer real-yield series dating from 1981, but is likewise too short. I use the *Market Volatility* data compilation associated with Shiller (1989) because it is the longest consistent annual time series that contains all of the variables required by the model in a single source that does not require additional splicing beyond the joins already documented above.

4.2 Yields and Bond Returns

This subsection details the construction of the real yield series and the resulting 10-year inflation-indexed bond returns. I use a parameter λ to construct alternative specifications of the 10-year yield, moving the financial environment from a baseline expectations-hypothesis case to a full Fisher-adjusted case. The procedure begins by generating an inflation forecast used to convert historical nominal yields into real terms. Let π_t denote the latest annual December-over-December inflation observable as of January of year t. Expected inflation follows the AR(1) process

$$\pi_{t+1} = a + \phi\pi_t + u_{t+1}.$$

Using these forecasts, I compute the real short rate and the Fisher-adjusted 10-year yield directly from the historical nominal yields as

$$\begin{aligned} y_{1,t} &= y^N_{1,t} - \mathbb{E}_t\pi_{t+1}, \\ y^{FA}_{10,t} &= y^N_{10,t} - \frac{1}{10}\sum_{j=1}^{10}\mathbb{E}_t\pi_{t+j}. \end{aligned}$$

This step relies on two identifying assumptions. First, the AR(1) specification represents the market's assessment of expected inflation. Second, the empirical construction makes no adjustment for an inflation risk premium. Under this baseline assumption, the Fisher-adjusted 10-year yield, $y^{FA}_{10,t}$, acts as an empirical upper bound that subsumes both the average level and any time-series variation of historical inflation risk premia into the measured real term premium. To construct the corresponding expectations-hypothesis (EH) lower bound, I model the joint dynamics of the real short rate and the Fisher-adjusted term spread, $spr^{FA}_t = y^{FA}_{10,t} - y_{1,t}$. An auxiliary bivariate linear model captures this rate process,

$$\begin{pmatrix} y_{1,t+1} \\ spr^{FA}_{t+1} \end{pmatrix} = \mu + \Psi \begin{pmatrix} y_{1,t} \\ spr^{FA}_t \end{pmatrix} + \varepsilon_{t+1}.$$

This system generates the EH-implied 10-year yield, defined as the average expected future path of short rates,

$$y^{EH}_{10,t} = \frac{1}{10}\sum_{j=0}^{9}\mathbb{E}_t y_{1,t+j}.$$

The gap between the Fisher-adjusted upper bound and the EH-implied lower bound defines the estimated real term premium,

$$TP_t = y^{FA}_{10,t} - y^{EH}_{10,t}.$$

With the two bounds established, the parameter λ constructs the alternative 10-year real yields that govern bond returns, according to

$$y_{10,t}(\lambda) = y_{10,t}^{EH} + \lambda TP_t.$$

Setting $\lambda = 0$ strictly enforces the EH-implied lower bound, $\lambda = 1$ recovers the full Fisher-adjusted upper bound, and $\lambda = 0.5$ yields the midpoint. Each value of λ therefore generates a corresponding real term spread,

$$spr_t(\lambda) = y_{10,t}(\lambda) - y_{1,t}.$$

Consequently, λ alters both the long-bond return series and the set of financial state variables used to predict returns. The equity valuation predictor cape_t and the one-year real yield $y_{1,t}$ remain common across all specifications. In contrast, the term spread $spr_t(\lambda)$ and the resulting expected bond returns are λ-specific. While λ remains explicit during this construction, I drop the index and write simply spr_t for notational convenience once a return system is fixed.

The next step translates the λ-specific yield series into corresponding long real bond returns. Following Campbell et al. (2003), I apply the Campbell–Lo–MacKinlay log-return approximation, which substitutes the newly observed n-year yield for the next-period $(n-1)$-year yield in the aging-bond formula. Unlike their real-consol application, however, the traded risky asset in this model is a finite 10-year synthetic real bond with par-coupon Macaulay duration. I compute the duration of this bond using its λ-specific yield, $y_{10,t}(\lambda)$, as

$$D_t(\lambda) = \frac{1 - (1 + Y_{10,t}(\lambda))^{-10}}{1 - (1 + Y_{10,t}(\lambda))^{-1}}, \qquad Y_{10,t}(\lambda) = \exp\left(y_{10,t}(\lambda)\right) - 1.$$

The resulting approximate log return is

$$r_{t+1}^{B}(\lambda) = D_t(\lambda) y_{10,t}(\lambda) - \left(D_t(\lambda) - 1\right) y_{10,t+1}(\lambda).$$

The corresponding log excess return is

$$xb_{t+1}(\lambda) = r_{t+1}^{B}(\lambda) - y_{1,t}.$$

Table 1 reports sample moments and Sharpe ratios for the 1920–2011 period, with mean excess returns Jensen-corrected to approximate simple returns. Because λ alters only the term premium, the properties of equities and the real short rate are identical across all specifications: stocks generate a mean excess return of 6.7% with 19.3% volatility (a Sharpe ratio of 0.346), and the short rate averages 2.0%. At $\lambda = 0$, the mean term spread is exactly zero because the unconditional expectation of future short rates equals the stationary mean of the short rate itself. The corresponding mean excess bond return rounds to zero, though the asset retains a return volatility of 5.9%. At $\lambda = 0.5$ and $\lambda = 1$, the mean term spread rises, directly lifting the mean excess bond return to 0.4% and 0.8%, respectively. Bond return volatility stays near 6% across these cases, reaching 6.3% at $\lambda = 1$. Even at this Fisher-adjusted upper bound, the long-bond Sharpe ratio of 0.129 remains well below that of equities.

Table 1: Sample Statistics

Sample moment	$\lambda = 0$	$\lambda = 0.5$	$\lambda = 1$
(1) $E[\text{xr}] + \sigma^2(\text{xr})/2$	0.067	0.067	0.067
(2) $\sigma(\text{xr})$	0.193	0.193	0.193
(3) SR = (1)/(2)	0.346	0.346	0.346
(4) $E[\text{xb}] + \sigma^2(\text{xb})/2$	0.000	0.004	0.008
(5) $\sigma(\text{xb})$	0.059	0.058	0.063
(6) SR = (4)/(5)	0.005	0.071	0.129
(7) $E[y_1]$	0.020	0.020	0.020
(8) $\sigma(y_1)$	0.029	0.029	0.029
(9) $E[\text{cape}]$	−2.727	−2.727	−2.727
(10) $\sigma(\text{cape})$	0.436	0.436	0.436
(11) $E[\text{spr}]$	0.000	0.004	0.007
(12) $\sigma(\text{spr})$	0.015	0.016	0.018

Notes: Data are from Shiller (1989). N = 92 annual observations (1920–2011). All variables in log units. xr is the excess log stock return. xb is the excess log long-bond return (both net of the log short rate). y_1 is the log short yield and cape is the CAPE-based log earnings yield, $-\log(\text{CAPE})$. spr denotes the λ-specific real term spread, $spr_t(\lambda) = y_{10,t}(\lambda) - y_{1,t}$. The variables xr, cape, and y_1 are common across λ; spr and xb are constructed separately for each λ. Sharpe ratios are Jensen-adjusted: SR = $(E[x] + \frac{1}{2}\hat{\sigma}_x^2)/\hat{\sigma}_x$. By construction, $E[\text{spr}] = 0$ at $\lambda = 0$.

4.3 VAR Specification and Inference

The return-state dynamics are modeled as a VAR(1) using annual data. A separate system is estimated for each parameter $\lambda \in \{0, 0.5, 1\}$. The financial state vector is given by

$$\mathbf{s}_t(\lambda) = (\text{cape}_t, \text{spr}_t(\lambda), y_{1,t})'$$

and the corresponding excess-return vector is

$$\mathbf{x}_{t+1}(\lambda) = (xr_{t+1}, xb_{t+1}(\lambda))' \,.$$

The stock excess return xr_{t+1}, the log equity valuation ratio cape_t, and the one-year real short rate $y_{1,t}$ are invariant across specifications. The real term spread $spr_t(\lambda)$ and the long-bond excess return $xb_{t+1}(\lambda)$ are reconstructed from the 10-year yield $y_{10,t}(\lambda)$. Varying λ therefore alters the state vector, predictive coefficients, and innovation covariance matrix alongside the mean bond return.

The system is estimated using the restricted mean-pinned OLS estimator of Campbell et al. (2003). To strictly enforce the expectations hypothesis at $\lambda = 0$, the VAR is additionally restricted to yield a zero conditional expected excess bond return. This sets the mean bond excess return and its forecasting coefficients to zero, as $\bar{x}_b = 0$ and $\Phi_{xb,\cdot} = 0$. The constraint eliminates the small predictable component ($R^2 = 1.5\%$) remaining in the raw $\lambda = 0$ returns due to the constant-maturity approximation.

To keep the state space of the life-cycle problem tractable, the VAR excludes lagged realized returns from the forecasting equations. Standard Granger-causality F-tests support this restriction: in the $\lambda = 1$ system, lagged xr and xb fail to predict cape, spr, and y_1 ($p = 0.40$, 0.46, and 0.96, respectively). An auxiliary unrestricted VAR also confirms no cross-return predictability or equity-return autocorrelation. The only significant return lag is a negative bond-return own-lag at $\lambda = 1$ ($\hat{\beta}_{xb} = -0.255$, $p = 0.017$), alongside a borderline joint test in the bond equation ($p = 0.052$). Given the sample size ($T = 92$) and the multi-stage construction of the real yields, this coefficient likely reflects finite-sample residual mean reversion. Furthermore, incorporating it would expand the financial state vector from three to five variables, making the dynamic program computationally intractable. Because the analysis specifically targets yield-curve predictability rather than generic return autocorrelation, I maintain the parsimonious state-based VAR.

To compute confidence intervals for the VAR estimates, I use a nonparametric stationary block bootstrap of an augmented annual donor panel following Politis and Romano (1994), drawing $B = 1000$ samples with a mean block length of eight years. A standard residual or transition-tuple bootstrap is invalid here because the variables entering the final VAR, particularly $y_{1,t}$, $\text{spr}_t(\lambda)$, and $xb_{t+1}(\lambda)$, are generated through a multi-stage empirical construction. The procedure instead resamples a pseudo-time sequence of annual donor rows over the 1920–2011 sample. Within each draw, the identical resampled index sequence is used across all λ specifications to ensure policy comparisons are paired on the same pseudo-history. Maintaining the inflation AR(1) parameters at their full-sample estimates, the auxiliary two-variable real-yield VAR is re-estimated on the pseudo-sequence to produce a draw-specific decomposition of the expectations-hypothesis long yield and the residual term premium.

A mechanical boundary issue arises because calculating annual bond and stock excess returns requires predecessor-year information. To address this, each donor row is augmented prior to resampling with its own historical predecessor inputs, including lagged nominal yields and inflation. Consequently, if a newly sampled block begins with historical year u succeeding an unrelated pseudo-predecessor year v, the return assigned to row u is computed exclusively from its explicitly attached $u - 1$ data. This prevents artificial holding-period returns across block joins while preserving valid historical one-year returns under the draw-specific λ construction. When the final VAR is estimated on the concatenated pseudo-series, the return in pseudo-row t is regressed on the state vector from pseudo-row $t - 1$. At a block join, this inevitably breaks the predictive transition between the lagged state and the contemporaneous return. This discontinuity is an artifact of how the block bootstrap approximates sampling variation in the dynamic system. Finally, conditional-mean restrictions such as imposing $\bar{x}_b = 0$ and $\Phi_{xb,\cdot} = 0$ at $\lambda = 0$ are enforced before the final VAR is estimated on each pseudo-series.

4.4 VAR Estimates and Intertemporal Hedging Demand

Table 2 reports the estimated VAR predictive coefficients, residual correlations, and bootstrap confidence intervals across the three term-premium specifications. In a dynamic portfolio choice

Table 2: VAR(1) Estimates by Term Premium Intensity (λ)

	Predictor variables					
	$\mathrm{cape_t}$	$\mathrm{spr_t}$	$\mathrm{y\_1_t}$	$\mathrm{xr_t}$	$\mathrm{xb_t}$	R^2
			Panel: $\lambda = 0$			
$\mathrm{cape_{t+1}}$	0.876* [0.555, 0.850]	−1.293 [−55.602, 37.395]	−0.363 [−41.910, 28.968]	–	–	0.800
$\mathrm{spr_{t+1}}$	−0.002 [−0.009, 0.006]	0.466 [−1.353, 0.412]	−0.108* [−1.372, −0.129]	–	–	0.451
$\mathrm{y\_1_{t+1}}$	0.004 [−0.007, 0.012]	0.620* [0.352, 2.490]	1.075* [0.855, 2.391]	–	–	0.649
$\mathrm{xr_{t+1}}$	0.116* [0.031, 0.214]	0.364 [−22.905, 32.763]	−1.185 [−18.444, 22.593]	–	–	0.100
$\mathrm{xb_{t+1}}$	–	–	–	–	–	–
max $\lvert\mathrm{eig}(\Phi)\rvert = 0.930$						
Residual correlations						
cape	1.00					
spr	0.14	1.00				
y_1	−0.09	−0.97	1.00			
xr	−0.98	−0.15	0.10	1.00		
xb	0.02	0.81	−0.93	−0.03	1.00	
			Panel: $\lambda = 0.5$			
$\mathrm{cape_{t+1}}$	0.876* [0.555, 0.850]	−0.738 [−7.473, 6.112]	−0.072 [−5.460, 3.343]	–	–	0.800
$\mathrm{spr_{t+1}}$	−0.003 [−0.009, 0.006]	0.606* [0.062, 0.667]	−0.004 [−0.292, 0.114]	–	–	0.373
$\mathrm{y\_1_{t+1}}$	0.004 [−0.007, 0.012]	0.354 [−0.036, 0.736]	0.936* [0.546, 1.075]	–	–	0.649
$\mathrm{xr_{t+1}}$	0.116* [0.031, 0.214]	0.208 [−3.628, 4.762]	−1.267 [−3.550, 1.786]	–	–	0.100
$\mathrm{xb_{t+1}}$	−0.013 [−0.028, 0.015]	1.117 [−0.213, 2.484]	0.449 [−0.385, 1.391]	–	–	0.046
max $\lvert\mathrm{eig}(\Phi)\rvert = 0.930$						
Residual correlations						
cape	1.00					
spr	0.16	1.00				
y_1	−0.09	−0.92	1.00			
xr	−0.98	−0.17	0.10	1.00		
xb	−0.04	0.49	−0.79	0.03	1.00	
			Panel: $\lambda = 1$			
$\mathrm{cape_{t+1}}$	0.876* [0.555, 0.850]	−0.516 [−4.327, 3.640]	0.044 [−3.458, 2.001]	–	–	0.800
$\mathrm{spr_{t+1}}$	−0.003 [−0.010, 0.007]	0.662* [0.246, 0.695]	0.056 [−0.155, 0.182]	–	–	0.374
$\mathrm{y\_1_{t+1}}$	0.004 [−0.007, 0.012]	0.248 [−0.018, 0.494]	0.880* [0.518, 0.906]	–	–	0.649
$\mathrm{xr_{t+1}}$	0.116* [0.031, 0.214]	0.145 [−2.103, 2.808]	−1.300 [−2.728, 0.624]	–	–	0.100
$\mathrm{xb_{t+1}}$	−0.010 [−0.036, 0.018]	1.505* [0.355, 2.634]	0.406 [−0.167, 1.130]	–	–	0.129
max $\lvert\mathrm{eig}(\Phi)\rvert = 0.930$						
Residual correlations						
cape	1.00					
spr	0.18	1.00				
y_1	−0.09	−0.88	1.00			
xr	−0.98	−0.18	0.10	1.00		
xb	−0.09	0.12	−0.58	0.08	1.00	

Notes: Point estimates with 90% bootstrap percentile interval [p_{05}, p_{95}] in square brackets. Here cape is the CAPE-based log earnings yield, $-\log(\text{CAPE})$. In each panel, spr and xb use that panel's λ-specific long yield. Stationary block bootstrap of Politis and Romano (1994), applied to the annual primitive panel, mean block length 8, paired across λ. $B = 1000$ draws under the headline (conditional AR(1)) specification. Estimation sample: 1920–2011 ($T = 92$). * marks coefficients whose 90% interval excludes zero on the same side as the point estimate. max $\lvert\mathrm{eig}(\Phi)\rvert$ is the largest absolute eigenvalue of Φ. Stationarity requires < 1. Dashed entries (–) are zero by restriction: lagged-return columns ($\mathrm{xr_t}$, $\mathrm{xb_t}$) are excluded from every equation. At $\lambda = 0$ the xb equation is restricted to zero under the maintained EH benchmark ($\mathbb{E}_t[\mathrm{xb_{t+1}}] = 0$ in the return–state system). For persistent state coefficients, bootstrap percentile intervals may lie below the full-sample estimate because block joins attenuate estimated persistence in finite samples.

model, a state variable induces intertemporal hedging demand if it is persistent, forecasts future expected returns, and its innovations covary with contemporaneous asset payoffs. The estimates confirm that all three state variables ($y_{1,t}$, spr_t, and cape_t) are highly persistent. Because the equity specification remains constant across all three λ regimes, it provides a straightforward illustration of these conditions. The valuation state cape_t positively and significantly predicts future stock excess returns ($\Phi_{xr,\text{cape}} = 0.116$), and its innovations correlate negatively with contemporaneous stock returns (-0.98). Consequently, poor current stock returns coincide with cheaper valuations and higher expected future returns, allowing stocks to hedge variations in their own risk premium. This standard valuation-based hedging motive (Campbell et al., 2003) operates identically across the different yield constructions and therefore cannot drive the main comparative statics.

Long bonds offer two intertemporal hedging channels. They can hedge future real short rates ($y_{1,t}$) or their own future excess returns (xb_{t+1}). Because the baseline $\lambda = 0$ specification restricts expected bond excess returns to zero, it eliminates the risk-premium channel. The asset functions solely as a real-rate hedge, generating positive returns when falling short rates worsen future reinvestment opportunities. The covariance estimates quantify this mechanism. By definition, the 10-year yield is $y_{10,t} = y_{1,t} + \text{spr}_t$. Table 2 reports that innovations to the short rate and the term spread strongly offset, sharing a residual correlation of -0.97. A decline in the short rate nonetheless lowers the overall 10-year yield. Due to the mechanical duration effect, bond returns inversely track these continuous long-yield innovations, with a correlation of -0.99. This structure ensures that a negative shock to $y_{1,t}$ drives a positive return in xb_{t+1}. Therefore, while bond returns correlate with both short-rate and spread innovations individually, the asset functions strictly to lock in a 10-year real rate.

Increasing λ to 0.5 and 1 activates the risk-premium channel. A bond held to maturity still locks in the 10-year real yield across all specifications. However, including the empirical term premium alters the underlying drivers of one-period excess bond returns (xb_{t+1}). Expected bond excess returns become time-varying, and at $\lambda = 1$, the term spread spr_t becomes a statistically significant predictor of future bond excess returns ($\Phi_{xb,\text{spr}} = 1.505$). While this predictive relationship is economically significant—at $\lambda = 1$, a one-standard-deviation increase in spr_t of 1.8 percentage points raises next-year expected log bond excess returns by approximately 2.8 percentage points—the hedging demand relies on the joint yield-curve system. Adding the term premium also changes the innovation correlations. Because the spread absorbs the volatile empirical term premium, the residual correlation between short-rate and spread innovations weakens from -0.97 at $\lambda = 0$ to -0.88 at $\lambda = 1$. Consequently, aggregate one-period bond returns no longer reflect pure revisions to expected future short rates. The correlation between bond excess returns and the short rate falls to -0.58, while its correlation with the spread shifts to 0.12. Taken together, the yield-curve innovations have the sign needed for intertemporal hedging: the negative covariance between bond returns and short-rate innovations, combined with the short rate positively predicting future bond excess returns, dominates the same-direction spread component.

The VAR estimates demonstrate weak cross-asset predictability and correlation. The short rate $y_{1,t}$ produces a large negative point estimate for forecasting stock excess returns xr_{t+1}, but the estimate is highly imprecise. At $\lambda = 1$, the short-rate coefficient is $\Phi_{xr,y_1} = -1.300$, so a one-standard-deviation increase in $y_{1,t}$ of 2.9 percentage points would lower next-year expected log equity excess returns by about 3.8 percentage points. However, the 90% bootstrap interval maps this effect to a wide range stretching from roughly -8.0 to +1.8 percentage points. Contemporaneous residual correlations between stock returns and bond returns xb_{t+1} are similarly weak, remaining near zero (-0.03 to 0.08) across specifications. The residual correlations between stock returns and innovations to the yield-curve variables ($y_{1,t}$ and spr_t) also remain close to zero. As documented by Gormsen and Lazarus (2026), interest-rate movements combine discount-rate and cash-flow news, this offset mutes the unconditional empirical relationship between equities and the yield curve.

Table 2 documents substantial estimation uncertainty in the VAR predictive equations. At $\lambda = 1$, the coefficient of the term spread predicting future bond returns is 1.505, but its 90% bootstrap interval ranges widely from 0.355 to 2.634, and the short-rate predictive coefficient spans zero. This imprecision is particularly consequential in a life-cycle setting. Because the financial state variables are highly persistent, small variations in these one-period forecasting coefficients compound over time, potentially driving large differences in long-horizon expected returns and the resulting intertemporal hedging demand. The subsequent portfolio choice model operates under the assumption that the estimated return system is known, relying entirely on the point estimates. Evaluating the full uncertainty bounds of the optimal allocations would require re-solving the dynamic portfolio problem for each block-bootstrap draw. Given the size of the state space, this procedure is computationally infeasible using the Endogenous Grid Method of Carroll (2006). Therefore, treating the point estimates as fixed, following Campbell et al. (2003), is necessary to tractably isolate the comparative statics of the term premium.

5 Results

This section presents the life-cycle portfolio implications of introducing term premia into the 10-year real bond. To evaluate this mechanism, the return specification shifts from the expectations-hypothesis baseline, which assumes a zero term premium, to the Fisher-adjusted 10-year yield, which contains the full estimated real term premium. The analysis first examines the expectations-hypothesis baseline ($\lambda = 0$) to establish the model's basic life-cycle properties. This isolates standard life-cycle forces from the bond-return predictability introduced by the estimated real term premium. Household preferences and background risks incorporate the core life-cycle features from Catherine et al. (2024): risky labor income, Social Security, mortality, and a bequest motive. The model expands this framework by incorporating time-varying risk premia and allowing wealth allocation across equities, 10-year real bonds, and short real bills. These baseline profiles demonstrate the model's standard life-cycle portfolio behavior and provide the reference point for the comparative statics across the λ scaling parameter.

5.1 Baseline Life-Cycle Allocation under the Expectations Hypothesis

Figure 1 reports simulated wealth, income, and portfolio shares for $\gamma = 5$ under the expectations-hypothesis baseline. The lines are medians among surviving households at each age, and the bands report the interquartile range. Labor income is strongly right-skewed during working life. At age 45, median labor income is 0.53 times average wage income, while mean labor income is 2.36. The interquartile range is 0.15 to 1.68. The Social Security formula compresses this dispersion sharply after retirement. At age 67, the interquartile range for benefits is 0.13 to 0.55, and the mean benefit is 0.32. Wealth is even more skewed. Median wealth rises from 0.57 at age 25 to 18.20 at retirement and 24.05 at age 80, while mean wealth at retirement is 87.71. The model therefore generates large upper-tail balance sheets and little retirement decumulation. This is not only a survivor-composition effect: among households alive at both ages 67 and 80, median wealth rises

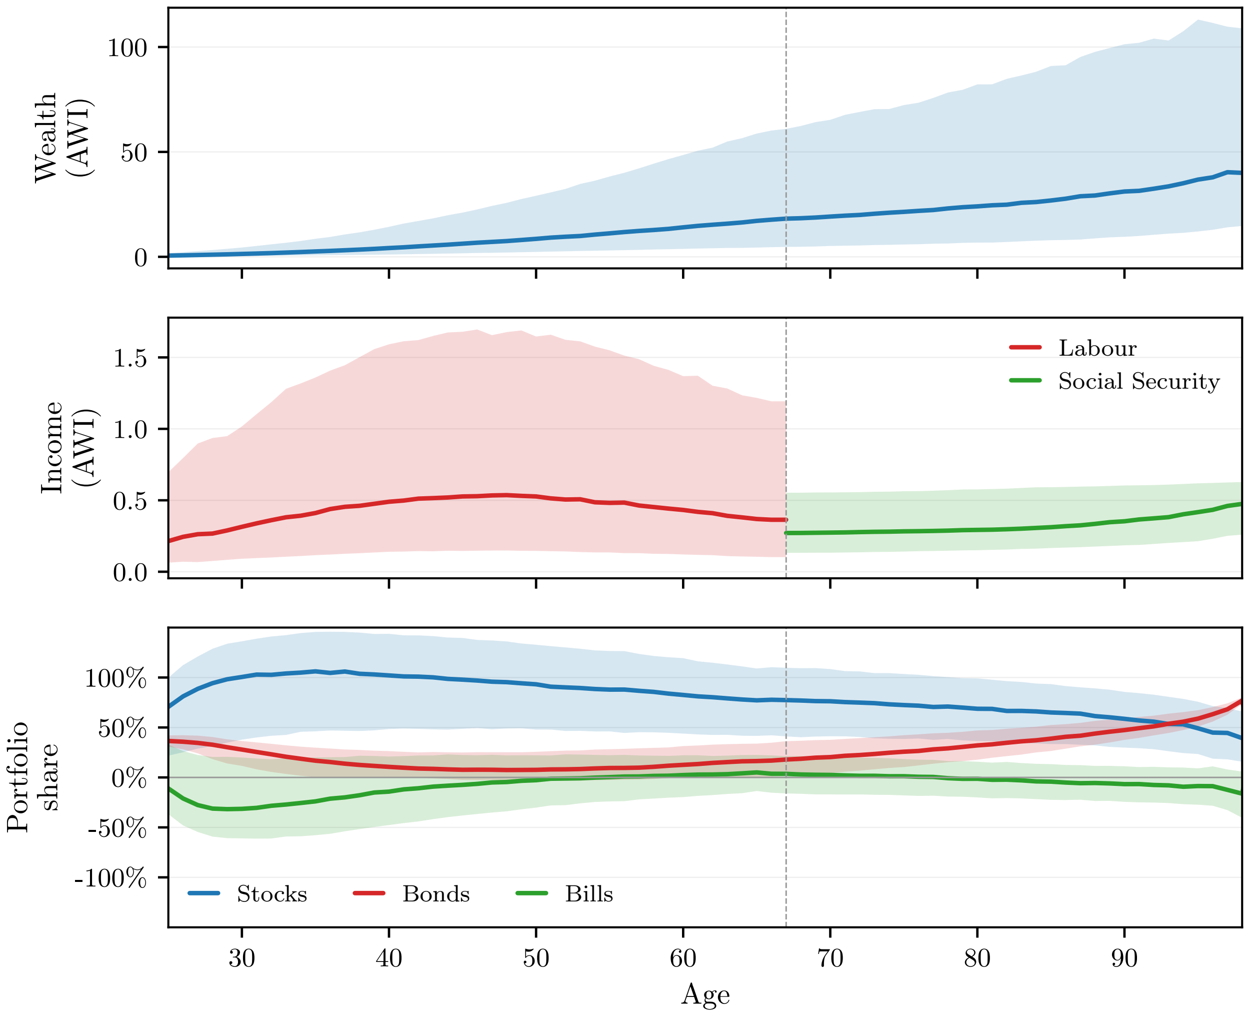

Figure 1: Cross-sectional medians and interquartile bands of wealth, income, and portfolio shares at $\gamma = 5$, $\lambda = 0$.

Notes: Top: household wealth in AWI units. *Middle:* disposable labor income (red, ages 22–66) and Social Security benefit (green, ages 67–99). *Bottom:* portfolio shares in stocks (blue), bonds (red), and short-term bills (green). Solid lines are cross-sectional medians and shaded bands are 25th–75th percentile bands across $N = 10{,}000$ simulated life-cycles. The dashed vertical line marks the statutory retirement age (67). Cross-sectional moments at each age are computed across agents alive at that age. Late-life moments reflect survivor selection, since the model has earnings-related mortality. Initial wealth 0.1 AWI, persistent income drawn from its AR(1) stationary distribution, initial VAR state at the joint-grid center.

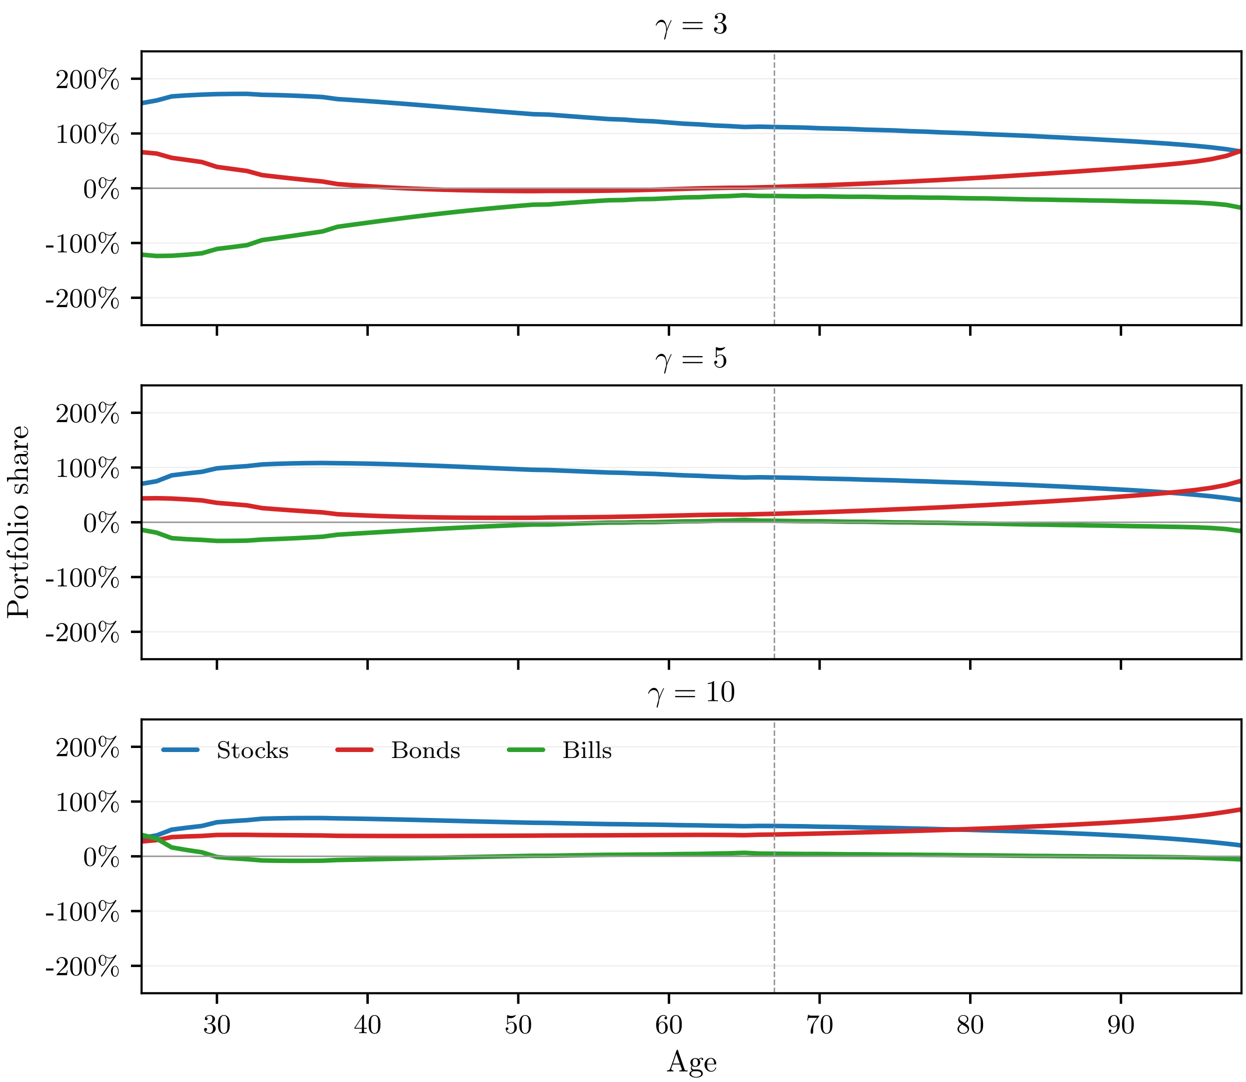

Figure 2: Policy portfolio shares by age for three levels of risk aversion $\gamma \in \{3, 5, 10\}$, all evaluated at the EH-coherent baseline $\lambda = 0$.

Notes: Each panel reports the optimal share in stocks (blue), long-term bonds (red), and short-term bills (green) implied by the life-cycle policy function evaluated at the joint mean of the financial-state distribution and the mean of the persistent income state, and at the median simulated wealth at each age (taken from a 10,000-agent forward simulation of the $\gamma = 5$ bundle, initial wealth 0.1 AWI). The dashed vertical line marks the statutory retirement age (67). As risk aversion increases, the optimal portfolio shifts out of leveraged stocks (with shorted bills as financing). Bond demand rises sharply in late retirement under all three calibrations.

by 1.39 average-wage units, and only 33 percent have lower wealth at age 80 than at retirement. Two forces help explain this pattern. First, the bequest motive gives households direct utility from preserving wealth, especially at older ages when mortality risk is high. Second, expected portfolio returns are high relative to the rate of time preference. In the Euler equation, this raises desired consumption growth and weakens the incentive to run down assets early in retirement.

The portfolio profiles have the standard life-cycle equity pattern but a less standard long-bond pattern. The median stock share rises from 71 percent at age 25 to about 106 percent at age 35, and then declines to 77 percent at retirement, 69 percent at age 80, and 39 percent by age 98. This is consistent with the human-capital mechanism of Bodie et al. (1992), Viceira (2001), and Cocco et al. (2005). Young households hold a large present value of future labor income relative to financial wealth. Since this human capital is close to bond-like in the calibration, a Merton-style target risky share of total wealth implies a high, sometimes levered, risky share of financial wealth. As human capital is depleted with age, the financial equity share falls. The long-bond share is instead U-shaped. It is 37 percent at age 25, falls to 8 percent at age 45, and then rises to 18 percent at retirement, 32 percent at age 80, and 77 percent by age 98. A natural interpretation is that late in life the bequest branch receives increasing weight as mortality rises. Since bequests are valued as a 10-year annuity, 10-year real bonds partly hedge the value of this annuity-like object. Near the terminal age, the problem becomes close to a one-period portfolio problem over bequest wealth, rather than a standard working-life life-cycle allocation.

Figure 2 shows that these baseline patterns move in the expected direction with risk aversion. The peak stock share falls from 172 percent when $\gamma = 3$ to 108 percent when $\gamma = 5$ and 70 percent when $\gamma = 10$. At age 45, the stock share falls from 149 percent to 103 percent and then to 65 percent as γ rises across the same values. Households substitute toward 10-year real bonds: the age-45 bond share rises from -3 percent to 8 percent and then to 37 percent. The short bill position also becomes less negative, its most negative value is -124 percent when $\gamma = 3$, compared with -34 percent when $\gamma = 5$ and -8 percent when $\gamma = 10$. Financial wealth therefore becomes less risky and less levered as risk aversion rises, while the qualitative life-cycle shapes remain similar across calibrations.

5.2 Term Premia and Life-Cycle Portfolio Choice

Figure 3 reports the optimal portfolio allocations for $\gamma = 5$ as the λ scaling parameter shifts the return specification from the expectations-hypothesis baseline ($\lambda = 0$) to the Fisher-adjusted 10-year yield ($\lambda = 1$). The policies are evaluated at the joint mean of the financial state and persistent-income state, and at the age-specific median wealth from the baseline ($\lambda = 0$) simulation. Applying this shared wealth path isolates true shifts in the underlying policy rules, preventing them from being confounded by endogenous wealth-dependent portfolio responses that would arise if each return regime generated its own evaluation points. The central finding is that introducing the estimated real term premium produces a large increase in 10-year real bond demand. At age 45, the allocation to this asset rises from 8 percent at $\lambda = 0$ to 33 percent at $\lambda = 0.5$ and 75 percent at $\lambda = 1$. At retirement, the corresponding shares are 16, 34, and 67 percent. Crucially, this increased bond allocation is not financed by crowding out equity. At age 45, the equity share falls modestly from 103 percent to 89 percent as λ rises from 0 to 1, while the short bill position drops from -12 percent to -64 percent. Consequently, shifting to the Fisher-adjusted 10-year yield expands the combined equity and 10-year real bond position from 111 percent to 164 percent of financial

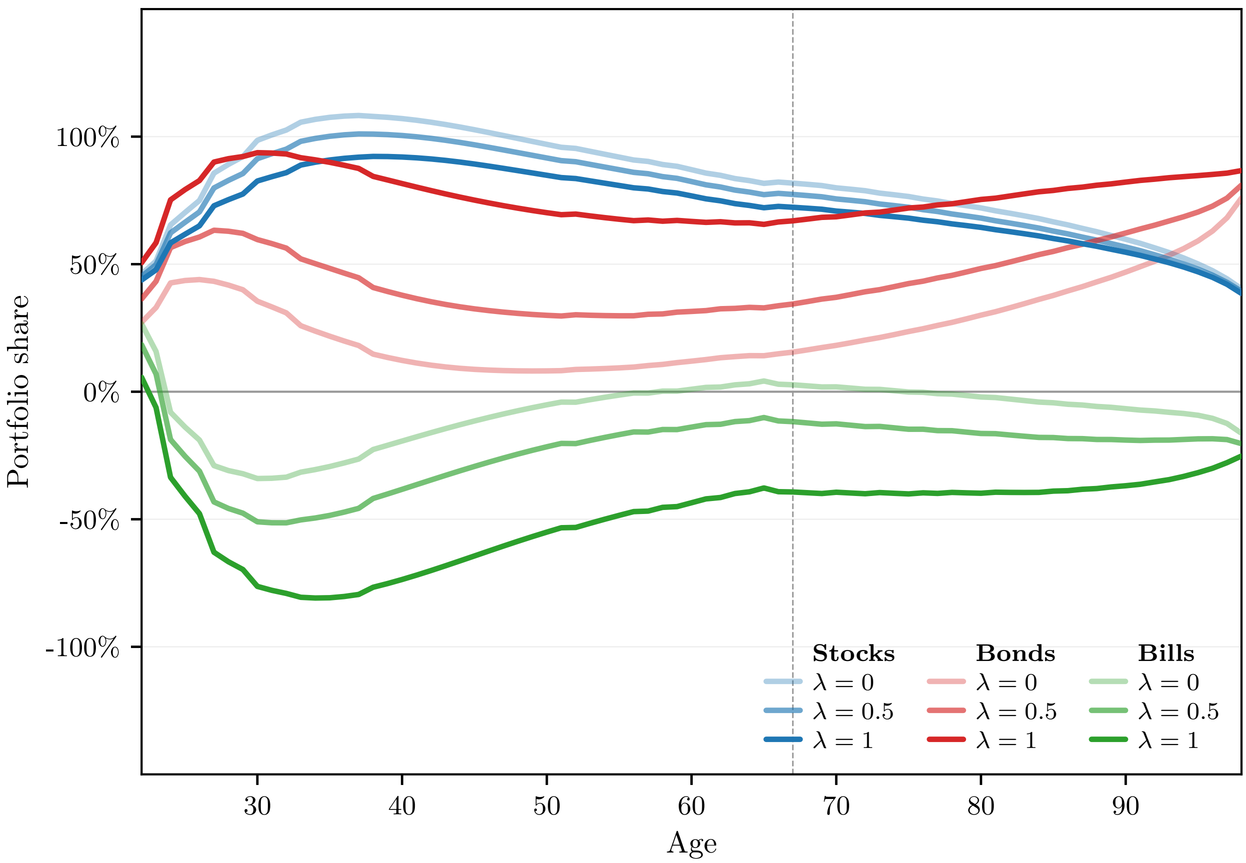

Figure 3: Life-cycle portfolio shares at $\gamma = 5$ for three term-premium intensities $\lambda \in \{0, 0.5, 1\}$.

Notes: Each line reports the optimal share in stocks (blue), 10-year bonds (red), or short-term bills (green) implied by the life-cycle policy function evaluated at the joint mean of the persistent income state z and the financial state $(\overline{\text{cape}}, \overline{\text{spr}}, \overline{y_1})$, and at the age-specific median simulated wealth path (from a 10,000-agent forward simulation of the $\gamma = 5$, $\lambda = 0$ bundle: initial wealth 0.1 AWI, persistent income drawn from its stationary distribution, initial VAR state at the joint-grid center, median wealth rises from 0.24 AWI at age 22 to roughly 37 AWI at age 95). Color intensity encodes λ: faint = $\lambda = 0$, vivid = $\lambda = 1$. The dashed vertical line marks the statutory retirement age (67).

wealth, financed predominantly by shorting short real bills. Section 6 discusses this unconstrained balance-sheet expansion and contrasts it with models that impose retail leverage constraints.

A portion of this increased demand can be viewed through a static mean-variance lens. As the scaling parameter λ rises from 0 to 1, the mean log excess return on the 10-year real bond increases from 0.0 percent to 0.61 percent, while its return volatility remains relatively stable near 6.3 percent. The Jensen-adjusted mean used in the Sharpe-ratio numerator rises to 0.8 percent, so the asset's Sharpe ratio moves from 0.005 to 0.129. In a standard one-period mean-variance portfolio with $\gamma = 5$, this changes the implied 10-year real bond share from -2.6 percent to 24.4 percent. Furthermore, despite the larger combined stock and bond position, the one-period return volatility of the portfolio does not rise. At age 45, the combined risky share rises from 111 percent to 164 percent as λ moves from 0 to 1, but the portfolio excess-return standard deviation actually falls from 19.8 percent to 18.3 percent. This reallocation reduces overall portfolio volatility because the decline in the equity share offsets the added 10-year real bond exposure, which is much less volatile than equities and maintains a near-zero correlation with stock returns. This mean-variance calculation confirms that while a higher Sharpe ratio increases static demand, it falls well short of explaining the full 67-percentage-point expansion. The remainder of the portfolio shift is instead driven by dynamic return predictability.

The remainder of the portfolio shift is driven by an intertemporal hedging channel. Returns on the 10-year real bond are mechanically linked to yield changes: an unexpected rise in the long yield inflicts an immediate capital loss, but simultaneously improves future reinvestment opportunities. Under the expectations-hypothesis baseline ($\lambda = 0$), expected 10-year real bond excess returns are restricted to zero, meaning this improvement operates exclusively through higher expected future short rates. However, under the Fisher-adjusted 10-year yield ($\lambda = 1$), the same joint yield-curve shock also raises expected future excess returns on the 10-year real bond. Because capital losses coincide with upward revisions to future bond risk premia through the combined short-rate and spread state, the 10-year real bond serves as an intertemporal hedge against variation in its own expected return. Numerically, the bond excess-return equation depends jointly on the term spread and the short rate, with predictive coefficients $\Phi_{xb,\text{spr}} = 1.505$ and $\Phi_{xb,y_1} = 0.406$, yielding an R^2 of 0.129. To quantify the magnitude of this channel, the dynamic policy is compared against an IID counterfactual that preserves the life-cycle environment and the unconditional return moments but removes all return predictability.

Table 3 decomposes the total portfolio allocation into a static IID component and a dynamic predictability component:

$$\alpha_t^{\text{full}}(\lambda) = \alpha_t^{\text{iid}}(\lambda) + \alpha_t^{\text{hedge}}(\lambda).$$

The IID model holds the unconditional return mean $\mu_r(\lambda)$ and covariance $\Sigma_{rr}(\lambda)$ fixed, but restricts the return-predictability block Φ_{xs} to zero. Additionally, the yield-curve variables used to price the bequest annuity are fixed at their unconditional sample means. The residual α^{hedge} therefore isolates intertemporal hedging demand, capturing the portion of the portfolio driven by time-varying expected returns.

Table 3 shows that the static IID component accounts for a minority of the total portfolio response. At age 45, the IID allocation to the 10-year real bond increases from 32 percent under the expectations-hypothesis baseline ($\lambda = 0$) to 57 percent under the Fisher-adjusted 10-year yield ($\lambda = 1$). This provides 25 percentage points of the total 67-percentage-point increase. The ma-

Table 3: IID and hedging demand by age, by term-premium intensity λ

Age	30	35	40	45	50	55	60	65	70	75	80	85	90
Panel: $\lambda = 0$													
IID demand													
Stocks	0.87	0.80	0.71	0.64	0.58	0.54	0.50	0.47	0.45	0.44	0.42	0.40	0.38
Bonds	0.44	0.40	0.36	0.32	0.29	0.27	0.25	0.24	0.23	0.22	0.21	0.20	0.19
Bills	−0.31	−0.21	−0.07	0.03	0.12	0.19	0.24	0.29	0.32	0.34	0.38	0.40	0.42
Hedging demand													
Stocks	0.12	0.27	0.36	0.38	0.39	0.38	0.37	0.35	0.35	0.33	0.31	0.27	0.22
Bonds	−0.08	−0.19	−0.24	−0.24	−0.21	−0.18	−0.13	−0.10	−0.05	0.02	0.09	0.18	0.28
Bills	−0.03	−0.09	−0.12	−0.15	−0.18	−0.20	−0.23	−0.25	−0.30	−0.35	−0.40	−0.45	−0.49
Panel: $\lambda = 0.5$													
IID demand													
Stocks	0.86	0.79	0.70	0.63	0.58	0.53	0.50	0.47	0.45	0.43	0.41	0.39	0.38
Bonds	0.56	0.52	0.46	0.41	0.38	0.35	0.33	0.30	0.29	0.28	0.27	0.26	0.25
Bills	−0.42	−0.31	−0.16	−0.05	0.05	0.12	0.18	0.23	0.26	0.29	0.32	0.35	0.38
Hedging demand													
Stocks	0.05	0.21	0.30	0.33	0.34	0.34	0.33	0.31	0.31	0.29	0.27	0.24	0.19
Bonds	0.03	−0.03	−0.08	−0.09	−0.08	−0.05	−0.01	0.02	0.08	0.14	0.22	0.29	0.38
Bills	−0.09	−0.17	−0.22	−0.24	−0.27	−0.29	−0.31	−0.33	−0.39	−0.44	−0.49	−0.53	−0.57
Panel: $\lambda = 1$													
IID demand													
Stocks	0.84	0.78	0.69	0.62	0.56	0.52	0.49	0.46	0.44	0.42	0.40	0.38	0.37
Bonds	0.77	0.71	0.63	0.57	0.52	0.48	0.45	0.42	0.40	0.39	0.37	0.35	0.34
Bills	−0.61	−0.49	−0.32	−0.19	−0.08	0.00	0.07	0.13	0.16	0.19	0.23	0.26	0.29
Hedging demand													
Stocks	−0.02	0.13	0.23	0.27	0.29	0.29	0.28	0.27	0.27	0.26	0.24	0.22	0.18
Bonds	0.17	0.19	0.19	0.18	0.19	0.20	0.22	0.24	0.28	0.33	0.39	0.44	0.48
Bills	−0.15	−0.32	−0.42	−0.46	−0.47	−0.49	−0.50	−0.50	−0.55	−0.59	−0.63	−0.65	−0.66

Notes: Policy portfolio shares evaluated at the joint mean of the persistent-income state z and the financial state $(\text{cape}, \text{spr}, y_1)$, where cape is the CAPE-based log earnings yield, $-\log(\text{CAPE})$, and at the age-specific median simulated wealth (from a 10,000-agent simulation with $\gamma = 5$, $\lambda = 0$). The IID demand is the policy from a life-cycle solve with i.i.d. returns (no VAR predictability), evaluated at the same point. Hedging demand is the residual: total demand minus IID demand. By construction, the three asset rows within each block sum to one (IID) and zero (hedging) at every age.

jority of the response is instead driven by the predictability component. Under the expectations-hypothesis baseline, intertemporal hedging demand for the 10-year real bond is negative, subtracting 24 percentage points from the midlife allocation. Introducing the estimated real term premium reverses this effect. At $\lambda = 1$, the hedging component becomes positive, contributing 18 percentage points to the 10-year real bond share. This 42-percentage-point shift in intertemporal hedging demand therefore constitutes roughly three-fifths of the overall increase. The hedging components for the remaining assets specify the financing margin. As λ shifts from 0 to 1, equity hedging demand remains positive but declines from 38 percentage points to 27 percentage points, while the bill hedging component drops from -15 percentage points to -46 percentage points. The dynamic expansion in 10-year real bond demand is thus accommodated predominantly through a larger short position in real bills.

Beyond shifting the average level of asset demand, return predictability dictates that optimal allocations fluctuate dynamically as investment opportunities change. Figure 4 illustrates this state

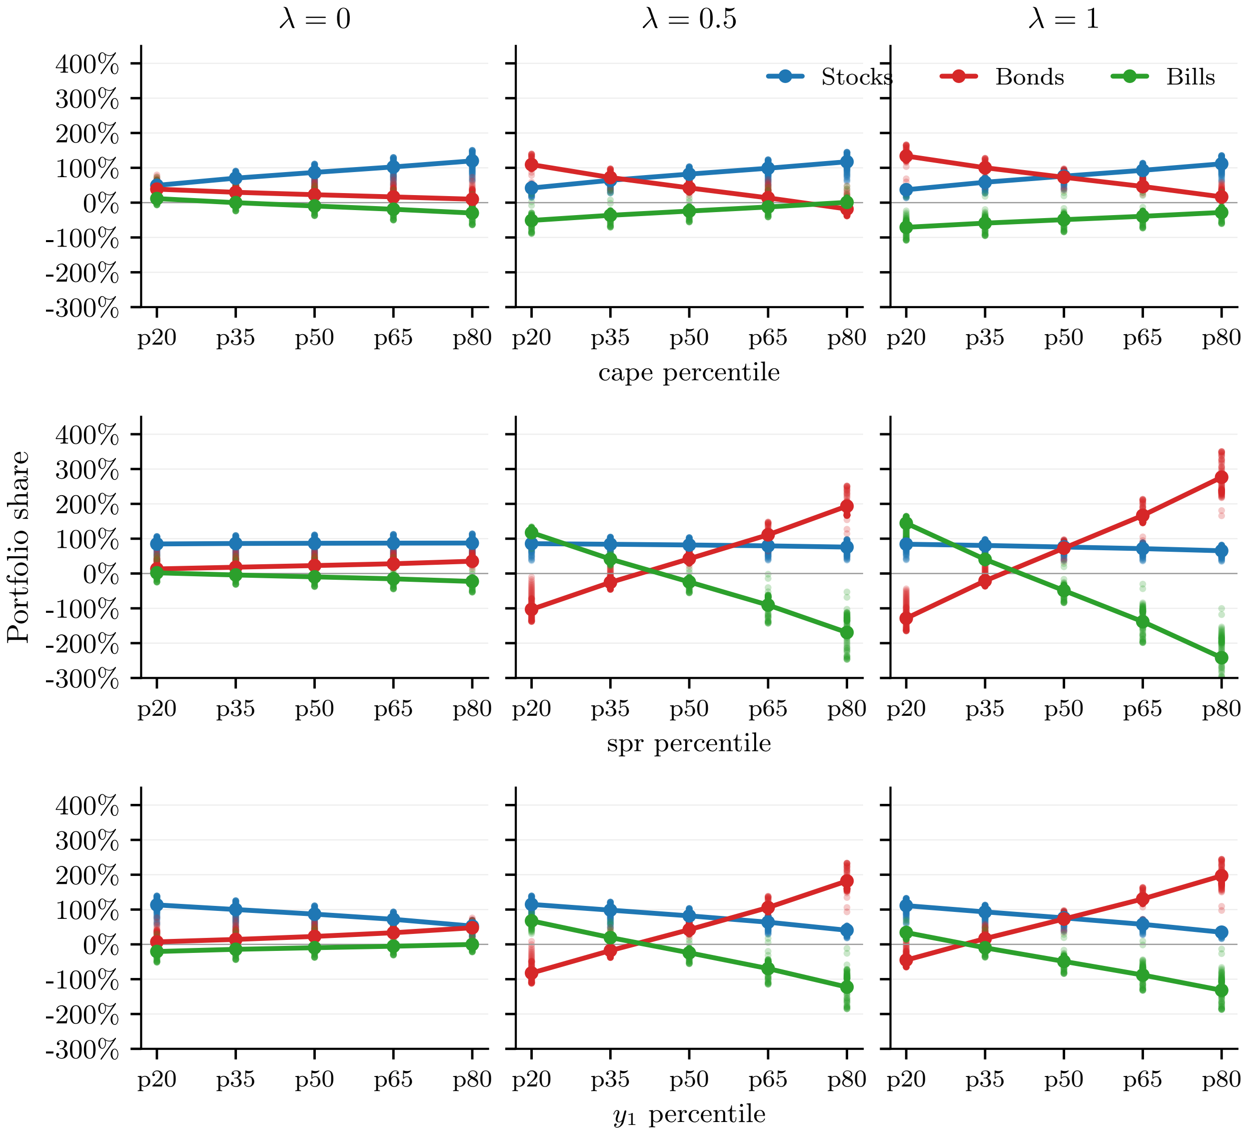

Figure 4: State-dependence of portfolio shares.

Notes: Each row varies one financial state variable through its marginal percentiles $\{0.20, 0.35, 0.50, 0.65, 0.80\}$ while the other two are held fixed at their unconditional means. Columns index the term-premium intensity $\lambda \in \{0, 0.5, 1\}$. Within each panel, light dots show the policy at every model age (22–99). Bold lines are survival-weighted averages over age. Hue identifies the asset class (red = bonds, blue = stocks, green = bills). Policy is evaluated at the median persistent-income state z and the age-specific median simulated wealth from a 10,000-agent simulation of the $\gamma = 5$, $\lambda = 0$. All bundles use $\gamma = 5$.

dependence by reporting how optimal portfolio shares respond as a single financial state variable sweeps from its 20th to its 80th marginal percentile, holding the remaining two variables at their unconditional means. The overarching pattern is that the household actively tilts the portfolio into an asset when the shifting state variable predicts higher future excess returns for that asset. For equities, this mechanism operates similarly across all return specifications: moving the valuation state (cape) from its 20th to its 80th percentile raises the equity share by 70 to 75 percentage points across all three columns. Under the Fisher-adjusted 10-year yield ($\lambda = 1$), for example, the equity share rises from 37 percent to 112 percent, while the 10-year real bond share falls from 134 percent to 17 percent.

The contrast is sharper for the yield-curve variables. Under the expectations-hypothesis baseline ($\lambda = 0$), the term spread has only a muted tactical effect on 10-year real bond demand: moving the spread from its 20th to its 80th percentile raises the bond share from 13 percent to 35 percent. Under the Fisher-adjusted 10-year yield ($\lambda = 1$), by contrast, the same percentile movement activates a much stronger bond-timing motive, raising 10-year real bond demand from -129 percent to 277 percent. Because the equity share falls only modestly, from 84 percent to 65 percent, the vast majority of this adjustment is absorbed by short real bills, whose share drops from 145 percent to -242 percent. The short-rate response reveals a similar logic. While a higher short rate implies a higher immediate safe return, households actually decrease their bill holdings: under $\lambda = 1$, moving y_1 from its 20th to its 80th percentile lowers the bill share from 34 percent to -132 percent, while 10-year real bond demand rises from -45 percent to 197 percent. This substitution occurs because the short rate simultaneously predicts higher future excess returns on the 10-year real bond, causing the active bond-timing motive to dominate the static appeal of the safe rate.

5.3 Which State Variables Matter?

The preceding analysis established that intertemporal hedging demand drives the majority of the expansion in long bonds when moving to the Fisher-adjusted 10-year yield. To isolate exactly which economic variables generate this hedging demand, Table 4 decomposes the aggregate shift across nested return-predictability systems. Let $\Delta_\lambda \alpha_{b,t} \equiv \alpha_{b,t}(1) - \alpha_{b,t}(0)$ denote the change in the optimal long-bond allocation at age t. Expressing the baseline decomposition from Table 3 in terms of these changes yields:

$$\Delta_\lambda \alpha_{b,t}^{\text{full}} = \Delta_\lambda \alpha_{b,t}^{\text{iid}} + \Delta_\lambda \alpha_{b,t}^{\text{hedge}}$$

To identify which predictors matter, the total hedging response can be further partitioned into the marginal contribution of adding the equity predictor (CAPE), and the subsequent contribution of restoring the yield-curve states (y_1 and spr). This yields the full decomposition:

$$\Delta_\lambda \alpha_{b,t}^{\text{full}} = \underbrace{\Delta_\lambda \alpha_{b,t}^{\text{iid}}}_{\text{IID baseline}} + \underbrace{\Delta_\lambda \alpha_{b,t}^{\text{hedge, cape}}}_{\text{CAPE increment}} + \underbrace{\Delta_\lambda \alpha_{b,t}^{\text{hedge, yc}}}_{\text{Yield-curve increment}}$$

Thus, the table asks whether the λ-driven bond reallocation appears when equity predictability is introduced, or only when the yield-curve states are restored. At age 45, the exact unrounded decomposition is $66.35 = 24.46 - 0.52 + 42.41$ percentage points. The CAPE increment is negligible, confirming that the yield-curve states account for the large remaining hedging response.

Table 4 confirms that the yield-curve states drive the large shift into long bonds as λ increases. To illustrate how this yield-curve hedging interacts with equity predictability in the baseline empirical model, Figure 5 evaluates the full life-cycle portfolio under $\lambda = 1$ across four nested predictability systems. Each panel isolates a specific asset and plots allocations under an IID baseline (no predictability), a system with only the short rate (y_1), a system adding the term spread ($y_1 + \text{spr}$), and the full system that also includes the equity valuation predictor (cape). The underlying return series and unconditional distributions remain identical across the four systems; what changes is the information set used to forecast returns, which dictates how variance is split between predictable drift and unforecastable innovations. Appendix B reports the corresponding restricted-predictor VAR estimates.

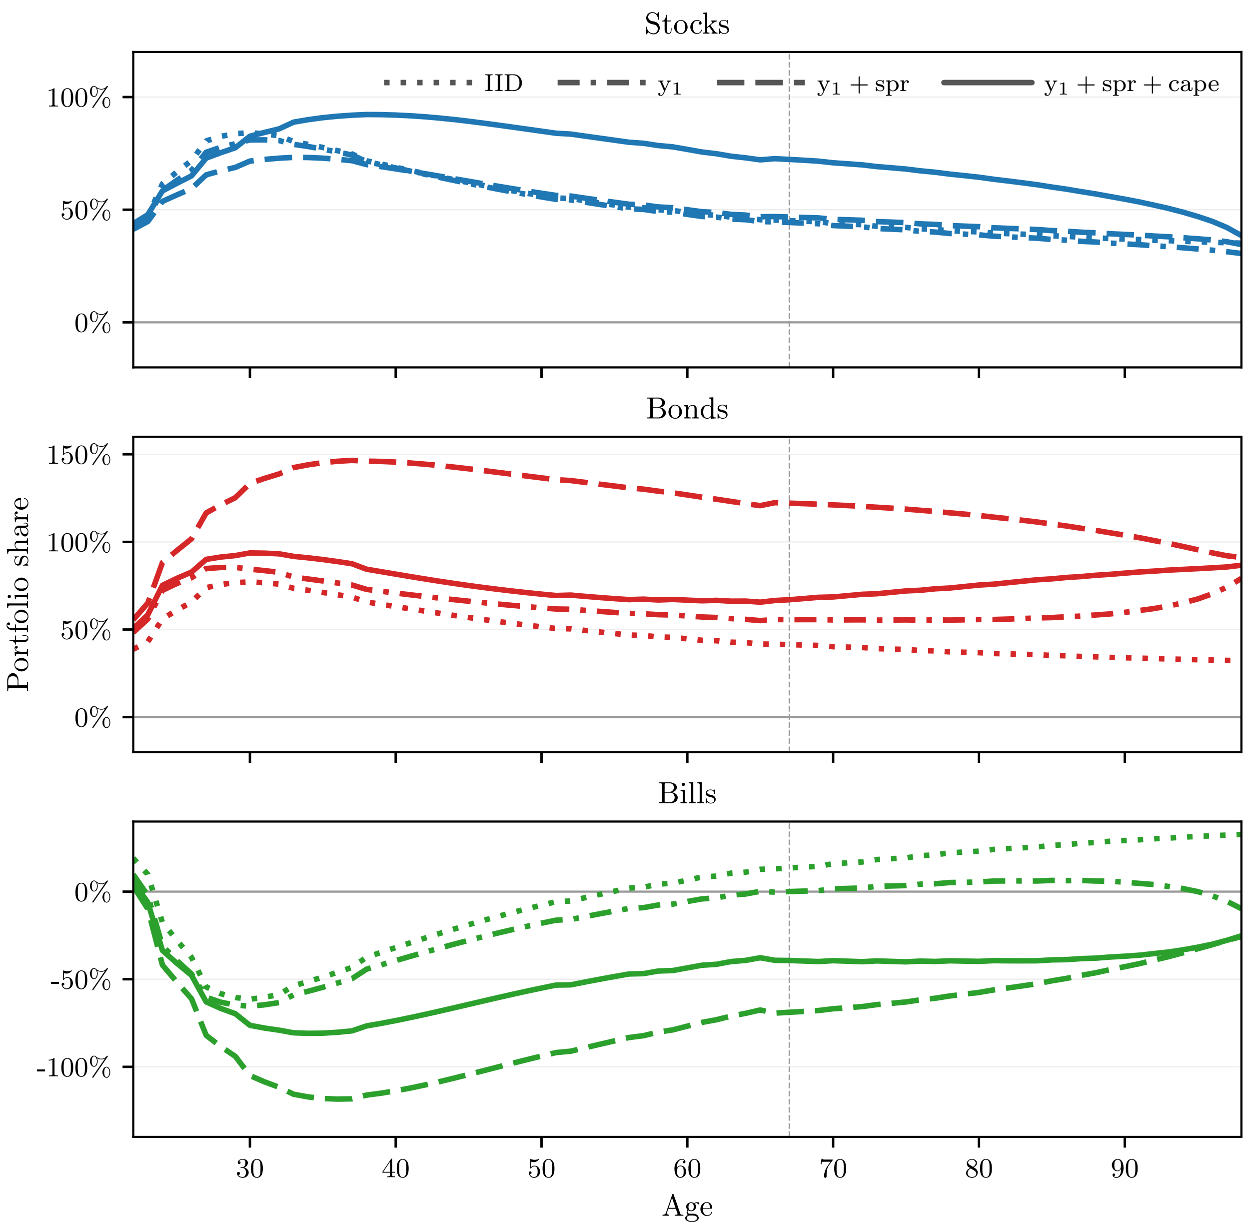

Figure 5: Life-cycle portfolio shares across four nested predictability systems.

Notes: One panel per asset (stocks, bonds, bills), each plotting four life-cycle profiles at $\gamma = 5$ and $\lambda = 1$. The four systems are nested by the state variables fed to the return process: IID (no predictability, dotted), y_1 alone (dashdot), $y_1 + \text{spr}$ (dashed), and the full $y_1 + \text{spr} + \text{cape}$ system (solid), where cape is the valuation predictor. For each system, policy is evaluated at the joint unconditional mean of the persistent-income state z and the financial state $(\text{cape}, \text{spr}, y_1)$ and at the age-specific median simulated wealth from a 10,000-agent simulation of the $\gamma = 5$, $\lambda = 0$. The vertical dashed line marks retirement at age 67.

Table 4: Sources of the λ-Driven Increase in Long-Bond Demand

Age	25	35	45	55	65	75	85
Total response: Full VAR	+36	+68	+66	+58	+51	+48	+41
Baseline: IID	+27	+31	+24	+21	+18	+17	+15
+ CAPE predictor	−2	−2	−1	0	0	0	+1
+ Yield-curve predictors	+11	+39	+42	+38	+34	+31	+25

Notes: Entries are percentage-point changes in the optimal long-bond allocation as the term-premium scaling parameter increases from $\lambda = 0$ to $\lambda = 1$. The IID model removes return predictability. The CAPE-only model retains the CAPE predictor while fixing the one-period bill return and bequest-annuity discount curve at their unconditional mean values. The Full VAR restores spr_t and $y_{1,t}$ as predictive states. Because the restricted VARs are re-estimated, the CAPE and yield-curve rows are model contributions, not pure hedging-demand terms. Components are calculated from unrounded policy values and may not sum exactly after rounding.

The nested profiles reveal a strict separation in hedging motives. The stock panel demonstrates that equity demand is driven almost entirely by the valuation predictor. In the IID system, the stock share peaks at 84 percent around age 30 before falling to 62 percent at age 45. Introducing the bond-side predictors (y_1 and spr) leaves this profile largely unchanged. It is only when cape is added in the full system that the stock share shifts upward, rising to 89 percent at age 45 and 64 percent at age 80. By the standard intertemporal hedging logic, stocks become more attractive to a long-horizon investor because negative return shocks reliably coincide with cheaper valuations and higher expected future returns.

Conversely, the bond panel shows that long-bond demand is driven by the yield-curve predictors. Adding y_1 raises bond holdings moderately, lifting the age-45 share from 57 percent in the IID system to 66 percent. However, introducing the term spread triggers a large expansion in demand. In the $y_1 + \text{spr}$ system, the age-45 bond share reaches 142 percent. This reflects the mechanical link between long-bond returns and yield-curve movements: because yield-curve variables strongly forecast future bond risk premia, capital losses are reliably offset by improved future investment opportunities.

Crucially, the full system shows that these two distinct hedging channels compete for the household's balance sheet. When cape is added to the yield-curve predictors, age-45 bond demand falls from 142 percent down to 75 percent, while equity demand expands. Thus, while the full model implies substantially higher bond demand than the IID benchmark, it generates much less bond demand than a system featuring only yield-curve predictability. Once the equity-timing channel becomes available, households optimally substitute a portion of their risk exposure away from long bonds to capture time-varying equity premia.

5.4 Infinite-Horizon Benchmark and Life-Cycle Effects

Table 5 compares the life-cycle allocation with an infinite-horizon benchmark facing the identical return process. The infinite-horizon benchmark strips the problem of life-cycle features, removing risky labor income, retirement, mortality risk, and the bequest motive. Mirroring our earlier separations of portfolio demand, the total life-cycle allocation at age t can be additively decomposed

Table 5: Life-cycle allocation relative to the infinite-horizon by term-premium intensity λ

Asset	α_{IH} (%)	Life-cycle deviation $\Delta = \alpha^{\text{lc}} - \alpha^{\text{ih}}$ (percentage points)						
		Δ age 25	Δ age 35	Δ age 45	Δ age 55	Δ age 65	Δ age 75	Δ age 85
Panel: $\lambda = 0$								
Stocks	75.7	−15.5	+14.2	+16.6	+10.7	−0.4	−3.5	−11.6
Bonds	12.8	+21.8	+5.6	−3.0	−1.9	+4.2	+12.5	+25.2
Bills	11.4	−6.3	−19.8	−13.6	−8.7	−3.8	−9.1	−13.7
Panel: $\lambda = 0.5$								
Stocks	71.9	−9.1	+13.1	+14.4	+9.1	−0.9	−4.1	−11.5
Bonds	24.1	+27.0	+13.9	+5.6	+5.2	+9.4	+18.0	+29.6
Bills	4.0	−18.0	−27.1	−20.0	−14.3	−8.5	−14.0	−18.1
Panel: $\lambda = 1$								
Stocks	67.0	+0.6	+11.9	+11.0	+6.6	−1.1	−3.8	−9.7
Bonds	53.7	+25.1	+15.5	+9.8	+8.0	+7.8	+13.9	+21.3
Bills	−20.6	−25.6	−27.3	−20.9	−14.6	−6.7	−10.1	−11.6

Notes: Life-cycle and infinite-horizon (IH) portfolio shares at $\gamma = 5$, evaluated at the unconditional mean of the financial state $(\text{cape}, \text{spr}, y_1)$, where cape is the CAPE-based log earnings yield, $-\log(\text{CAPE})$. The IH policy is wealth-invariant (no labor income, CRRA). The life-cycle policy is averaged across alive agents at each reference age, with each agent contributing the policy evaluated at their own persistent-income state z and wealth. $\Delta = \alpha_{\text{LC}} - \alpha_{\text{IH}}$. Both solves use the same VAR. The life-cycle simulation uses 10,000 agents with seed 42.

into a time-invariant infinite-horizon baseline and an age-dependent life-cycle deviation:

$$\alpha_t^{\text{lc}}(\lambda) = \underbrace{\alpha^{\text{ih}}(\lambda)}_{\text{IH baseline}} + \underbrace{\left(\alpha_t^{\text{lc}}(\lambda) - \alpha^{\text{ih}}(\lambda)\right)}_{\text{Life-cycle deviation}}$$

For each λ, the table reports both of these components. The residual life-cycle deviation isolates how the presence of human capital and finite-life frictions systematically distorts portfolio demand relative to the standard infinite-horizon baseline. The first observation is that the qualitative λ pattern is fundamentally a property of the return environment rather than the life-cycle structure. The baseline infinite-horizon bond share is relatively low at 12.8 percent under $\lambda = 0$. However, as more of the long-rate variation is explicitly attributed to term premia, this IH share expands to 24.1 percent at $\lambda = 0.5$ and 53.7 percent at $\lambda = 1$. Thus, the return process alone dictates the vast majority of the λ-driven increase in long-bond demand. The life-cycle features then systematically reshape the level and composition of this risk taking over time. In midlife, future human capital acts as a bond-like asset, allowing households to support higher financial risk. Consequently, the life-cycle stock-plus-bond position strictly exceeds the infinite-horizon baseline across every value of λ. At age 35, this excess life-cycle exposure is 19.8 percentage points at $\lambda = 0$ and expands to 27.3 percentage points at $\lambda = 1$. The composition of this excess risk also shifts with the return specification: at age 45 under $\lambda = 0$, households hold 16.6 percentage points more stocks than the infinite horizon, but slightly fewer bonds. When $\lambda = 1$ activates bond-return predictability, they expand both margins simultaneously, holding 11.0 percentage points more stock and 9.8 percentage points more bonds.

Finally, the composition of risk shifts late in life as mortality rises and the continuation value is increasingly driven by the bequest motive. The life-cycle adjustment becomes progressively bond-heavy. At $\lambda = 1$, excess bond demand relative to the infinite horizon rises from 7.8 percentage points at age 65 to 13.9 percentage points at age 75, reaching 21.3 percentage points at age 85. The identical late-life tilting appears at lower values of λ. Because the household increasingly prioritizes wealth preservation for the estate, it systematically substitutes away from equity volatility and into the long real bond. Ultimately, the empirical return process establishes the baseline level of asset demand, while life-cycle frictions govern how the overall size and composition of the household's risky portfolio evolve dynamically with age.

6 Discussion

6.1 Economic Interpretation of the Results

To understand the sharp increase in long-bond demand, it is necessary to unpack the return system. Moving from the expectations-hypothesis to the Fisher-adjusted 10-year yield does not merely introduce a constant term premium to the long bond's Sharpe ratio. Instead, this transition mechanically alters the time-series properties of the constructed 10-year yield and the corresponding term spread. Estimating the VAR on these redefined series alters both the predictive coefficients and the innovation covariance matrix, effectively inducing time-varying risk premia in bond returns. The resulting portfolio reallocation therefore captures two simultaneous effects: a static adjustment to higher unconditional risk compensation and a dynamic response to conditional expected bond returns.

Time-varying risk premia account for most of the portfolio response. At age 45, moving to the Fisher-adjusted real yield raises the long-bond allocation by 67 percentage points. A decomposition attributes 25 percentage points of this increase to static demand for higher unconditional bond returns, leaving 42 percentage points driven by intertemporal hedging. The mechanism relies on the joint yield-curve system forecasting future bond excess returns. When a rise in long yields generates an immediate bond capital loss, the associated increase in expected future bond compensation raises the continuation value. Long-bond risk is therefore less costly to the investor than its one-period volatility alone would suggest, because adverse current returns partly coincide with improved future investment opportunities. Finally, life-cycle features reshape the age profile of this demand. In mid-life, future labor income acts as an implicit safe asset, a role assumed by Social Security in retirement. Additionally, late-life mortality risk necessitates duration-matching the estate. These forces combine to elevate the household's total risk capacity above standard infinite-horizon benchmarks.

This expansion in long-bond demand is financed by shorting real bills rather than substituting for equities. At age 45, moving from the expectations-hypothesis baseline to the Fisher-adjusted yield reduces the equity share from 103 percent to 89 percent, while the short bill position drops from -12 percent to -64 percent. Because the long real bond is much less volatile than equity and remains only weakly correlated with stock returns, the leveraged bond expansion does not raise total portfolio risk; at age 45, annual portfolio volatility falls slightly, from 19.8 percent to 18.3 percent. The small decrease in the equity share offsets the variance of the added bond exposure. However, imposing borrowing constraints would prevent this balance-sheet expansion, forcing a strict substitution where holding more bonds requires holding less equity. Allowing unconstrained leverage reveals which margins the investor actively adjusts when free to choose its portfolio response. While this balance-sheet expansion may be infeasible for a retail household to implement directly, it aligns with the portfolio capabilities of delegated fund managers, institutional pension providers, and target-date funds with derivative overlays.

6.2 Relation to Term-Premium Portfolio-Choice Literature

The $\lambda = 0$ specification isolates the demand for a long real bond when it serves purely as a hedge against future real short rates. This expectations-hypothesis setting acts as a finite-maturity analogue to the real consol in Campbell et al. (2003). While both models generate comparable equity allocations, my infinite-horizon model yields a lower fixed-income demand: 12.8 percent in the 10-year bond compared to their 44.7 percent in the real consol. This disparity reflects differences in both the bond construction and preferences of the investor. Campbell et al. (2003) evaluate an investor with Epstein-Zin preferences who has access to a real consol, an asset that serves as the theoretical risk-free asset for infinitely lived investors. My setup models an investor with CRRA preferences and the asset menu instead includes a 10-year bond; the shorter maturity is potentially more appropriate for a life-cycle investor. Crucially, however, both their real consol and my $\lambda = 0$ baseline impose the expectations hypothesis. My results demonstrate that this shared restriction shuts down a source of bond demand: the ability to hedge against the bond's own time-varying risk premium.

Introducing the Fisher-adjusted yield at $\lambda = 1$ restores this risk-premium hedging channel and increases portfolio demand for bonds. This result aligns with Sangvinatsos and Wachter (2005), who demonstrate that time-varying bond risk premia generate large hedging demands for long-run investors. In their infinite-horizon consumption–portfolio problem with an affine term-structure model, optimal bond demand fluctuates widely across the values of the underlying state variables that govern expected excess bond returns through the time-varying price of risk. The state-dependent policies in my model capture a similar intuition: bond demand rises sharply when the observable term spread predicts higher future bond excess returns. However, the λ methodology evaluates a different problem. While Sangvinatsos and Wachter (2005) study variations in demand across states within a fixed term-structure model, shifting λ demonstrates how sensitive the optimal portfolio is to the empirical assumptions used to extract the term premium from historical yields.

The unconstrained portfolio response in my model provides a contrast to Koijen et al. (2010), the closest life-cycle comparison in the literature. Their model studies time-varying bond risk premia in a life-cycle setting with borrowing and short-sale constraints. State dependence in expected excess returns is concentrated on the fixed-income side: yield-curve factors drive time-varying bond premia, while the equity risk premium is held constant. Because early-life human capital and portfolio constraints limit the household's ability to adjust its fixed-income exposure, meaningful bond timing appears only around mid-life, when human capital has declined sufficiently. My unconstrained specification changes this margin. Rather than financing higher bond demand primarily by reducing equity, the household can expand the balance sheet by shorting bills, allowing it to capture the time-varying bond premium while maintaining substantial equity exposure. The comparison is further sharpened by the equity-predictability channel in my VAR. In the restricted-predictor exercise, removing the CAPE-based equity-timing channel raises bond demand relative to the full system, indicating that equity and bond predictability compete for portfolio weight on the household balance sheet.

6.3 Term-Premium Identification and Sensitivity

I construct real yields by subtracting a full-sample AR(1) inflation forecast from historical nominal yields. This approach provides a complete historical series, but it imposes the identifying assumption that the econometric forecast reflects contemporary market expectations. Estimating the AR(1) coefficients over the full 1920–2011 period also introduces a look-ahead bias into the historical inflation expectations. Crucially, the choice of inflation filter determines the properties of the measured real short rate, the 10-year term spread, and the expected bond returns evaluated in the VAR. Alternative adjustments mechanically alter the mean, persistence, and predictability of long-bond returns. For example, applying the seven-year moving-average adjustment used by Rogoff et al. (2024) would produce lower means and higher volatility for the 10-year real yield in this sample. Survey expectations offer an alternative that avoids econometric filtering, but they do not span the required time period. The Livingston Survey begins in 1946, omitting the first twenty-six years of data, and its early observations are difficult to compare against modern expected inflation measures. The AR(1) adjustment should therefore be understood as a parsimonious modeling choice rather than a unique estimate of real yields. Consequently, the quantitative bond-premium and hedging-demand results remain conditional on this chosen filter.

The $\lambda = 1$ specification uses the full Fisher-adjusted yield, attributing the entire gap between the empirical yield and the expectations-hypothesis baseline to the real term premium. However, because the Fisher adjustment subtracts only expected inflation from historical nominal yields, it implicitly assumes a zero inflation-risk premium. If investors demanded compensation for the risk that unexpected inflation would erode nominal payments, this inflation-risk compensation remains embedded in the constructed real yield. Defining the portfolio response strictly as demand for real-duration risk therefore overstates the mechanism if actual inflation risk comprises part of the scaled residual. Isolating the pure real premium would require estimating a joint term-structure model on inflation-indexed bonds. Because that modern sample is too short to estimate the persistent return dynamics required for the life-cycle problem, the paper instead relies on a sensitivity approach. Varying λ maps the portfolio response from a lower bound containing zero term premium to an upper bound that assigns the entire historical yield residual to the real bond, evaluating the maximal potential impact of term-risk compensation.

A final limitation concerns the statistical precision of the estimated return dynamics. The central dynamic mechanism relies on the term spread forecasting future bond returns, but estimating predictive regressions with persistent state variables over a short annual sample produces wide confidence intervals. For example, while the point estimate for the term spread forecasting the long-bond return at $\lambda = 1$ is 1.505, its 90% bootstrap interval is $[0.355, 2.634]$. Because repeatedly solving the high-dimensional life-cycle model across all bootstrap draws is computationally infeasible, the dynamic program treats the VAR point estimates as strictly given. Consequently, the parameter uncertainty from the return system is not propagated into the optimal portfolio rules. While the strictly positive bootstrap interval supports the qualitative presence of bond-return predictability, this estimation uncertainty means the exact quantitative magnitude of the intertemporal hedging demand is likely less robust than the fixed point estimates imply. This reinforces the underlying premise of the λ exercise: life-cycle portfolio allocations are highly sensitive to upstream econometric assumptions about the return-generating process.

6.4 Equity and Interest-Rate Risk

The analysis deliberately stops short of translating stock holdings into interest-rate exposure. The life-cycle model treats stocks and long real bonds as distinct assets with jointly predictable excess returns, and the λ sensitivity exercise focuses on how assumptions about long-bond returns affect bond demand. The model does not decompose stock returns into an interest-rate-sensitive component and a residual equity component. The reported bond allocation should therefore be interpreted strictly as demand for the traded long real bond, rather than a complete measure of the household's total interest-rate exposure.

This distinction matters in light of recent work on the interest-rate component of equity returns. Polk and Vuolteenaho (2026) decompose the conventional equity premium into a pure equity premium—defined as stock returns relative to duration-matched inflation-indexed bonds—and a real term-premium component. Their framework indicates that part of conventionally measured equity compensation reflects compensation for holding long-duration real claims rather than residual equity risk. Relatedly, van Binsbergen (2026) shows that benchmarking risky assets against duration-matched government bonds, rather than short bills, reclassifies part of the conventional excess return as compensation for duration exposure. Gormsen and Lazarus (2026) provide a related duration-based interpretation of equity factor premia, emphasizing the timing of firm cash flows. Taken together, this literature establishes that stocks contain economically important interest-rate exposure not separately identified in the present model.

Extending the life-cycle framework to decompose equity returns into an interest-rate component and a residual equity premium requires imposing additional structure on the baseline VAR. While the reduced-form VAR captures the unconditional covariance between stock returns and yield-curve innovations, this covariance conflates orthogonal interest-rate shocks with contemporaneous shifts in expected cash flows. Furthermore, calculating duration-adjusted equity returns requires explicit present-value restrictions to model the timing of future dividends. The present model leaves this structural decomposition aside. The portfolio results therefore capture how life-cycle households use the traded long real bond, leaving the broader question of total household interest-rate exposure across the entire balance sheet for future research.

7 Conclusion

This paper evaluates how assumptions about the real term premium alter optimal life-cycle portfolio decisions. To navigate the challenge of measuring true real interest rates over a long historical sample, I analyze the sensitivity of optimal portfolio choice to the scale and dynamics of the real term premium. By sweeping a scaling parameter, λ, I construct 10-year real yields that range from a zero-premium expectations-hypothesis lower bound up to a Fisher-adjusted upper bound. For each yield specification, I estimate an annual VAR to capture return predictability. I then embed these return dynamics into a structural life-cycle model featuring risky labor income, Social Security, mortality risk, and an annuity-equivalent bequest motive.

The central finding is that moving from the expectations-hypothesis baseline to the Fisher-adjusted 10-year yield drives a large expansion in the real bond's portfolio share, which rises from 8 percent to 75 percent at age 45. This increased bond allocation does not crowd out equities. Instead, households finance it by shorting short-term real bills. A decomposition reveals that this reallocation is driven primarily by intertemporal hedging rather than static demand for higher unconditional returns. Because capital losses coincide with upward revisions to future bond risk premia, the long bond serves as an intertemporal hedge against variation in its own expected return. Finally, while return dynamics dictate the overall scale of this bond demand, standard life-cycle forces such as human capital and Social Security shape its age profile.

These quantitative findings must be interpreted in light of the empirical challenges inherent in identifying the true real term premium. Estimating the persistent yield-curve dynamics required for a long-horizon portfolio problem necessitates a century-long historical sample, predating the existence of inflation-indexed bonds. Deflating historical nominal yields with an autoregressive inflation forecast implicitly assumes a zero inflation-risk premium. Consequently, the yield gap attributed to the real term premium in the Fisher-adjusted upper bound likely captures both the pure real term premium and historical compensation for inflation risk. The λ parameter addresses this ambiguity by transforming the estimation problem into a structured sensitivity analysis across the plausible bounds of the real term premium.

The primary contribution of this thesis lies in integrating the empirical VAR framework of Campbell, Chan, and Viceira (2003) into a structural life-cycle model. This approach demonstrates that optimal portfolio choice is highly sensitive to real term premium assumptions, despite the low Sharpe ratio of the long real bond. In a static mean-variance framework, this low average compensation would generate only minor demand. However, the dynamic aspects of portfolio choice alter this calculus. Because expected excess returns are predictable, the bond's intertemporal hedging value justifies a large portfolio share. Consequently, assumptions about the real term premium reshape life-cycle bond demand primarily through the return predictability they imply, rather than through the long real bond's modest unconditional compensation.

References

Ang, A., Bekaert, G., and Wei, M. (2008). The term structure of real rates and expected inflation. *Journal of Finance*, 63(2):797–849.

Barberis, N. (2000). Investing for the long run when returns are predictable. *Journal of Finance*, 55(1):225–264.

Bauer, M. D., Rudebusch, G. D., and Wu, J. C. (2012). Correcting estimation bias in dynamic term structure models. *Journal of Business and Economic Statistics*, 30(3):454–467.

Bodie, Z., Merton, R. C., and Samuelson, W. F. (1992). Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control*, 16(3–4):427–449.

Brennan, M. J., Schwartz, E. S., and Lagnado, R. (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control*, 21(8–9):1377–1403.

Brennan, M. J. and Xia, Y. (2002). Dynamic asset allocation under inflation. *Journal of Finance*, 57(3):1201–1238.

Campbell, J. Y., Chan, Y. L., and Viceira, L. M. (2003). A multivariate model of strategic asset allocation. *Journal of Financial Economics*, 67(1):41–80.

Campbell, J. Y. and Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird's eye view. *Review of Economic Studies*, 58(3):495–514.

Campbell, J. Y., Shiller, R. J., and Viceira, L. M. (2009). Understanding inflation-indexed bond markets. *Brookings Papers on Economic Activity*, 2009(1):79–120.

Campbell, J. Y. and Viceira, L. M. (1999). Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics*, 114(2):433–495.

Campbell, J. Y. and Viceira, L. M. (2001). Who should buy long-term bonds? *American Economic Review*, 91(1):99–127.

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

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

Catherine, S., Miller, M., Paron, J. D., and Sarin, N. (2024). Interest-rate risk and household portfolios. Jacobs Levy Equity Management Center for Quantitative Financial Research Paper.

Chetty, R., Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A., and Cutler, D. (2016). The association between income and life expectancy in the united states, 2001–2014. *Journal of the American Medical Association*, 315(16):1750–1766.

Cocco, J. F., Gomes, F. J., and Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies*, 18(2):491–533.

Cochrane, J. H. and Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, 95(1):138–160.

Duarte, V., Fonseca, J., Goodman, A. S., and Parker, J. A. (2021). Simple allocation rules and optimal portfolio choice over the lifecycle. NBER Working Paper No. 29559.

Fama, E. F. and Bliss, R. R. (1987). The information in long-maturity forward rates. *American Economic Review*, 77(4):680–692.

Friedman, M. and Schwartz, A. J. (1963). *A Monetary History of the United States, 1867–1960*. Princeton University Press, Princeton, NJ.

Gomes, F., Michaelides, A., and Zhang, Y. (2022). Tactical target date funds. *Management Science*, 68(4):3047–3070.

Gormsen, N. J. and Lazarus, E. (2026). Interest rates and equity valuations. NBER Working Paper No. 34814.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of u.s. workers reveal about lifecycle earnings dynamics? *Econometrica*, 89(5):2303–2339.

Guvenen, F., Ozkan, S., and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660.

Kim, T. S. and Omberg, E. (1996). Dynamic nonmyopic portfolio behavior. *Review of Financial Studies*, 9(1):141–161.

Koijen, R. S. J., Nijman, T. E., and Werker, B. J. M. (2010). When can life cycle investors benefit from time-varying bond risk premia? *Review of Financial Studies*, 23(2):741–780.

Lynch, A. W. and Tan, S. (2011). Labor income dynamics at business-cycle frequencies: Implications for portfolio choice. *Journal of Financial Economics*, 101(2):333–359.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1):77–91.

Meltzer, A. H. (2003). *A History of the Federal Reserve, Volume 1: 1913–1951*. University of Chicago Press, Chicago.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413.

Michaelides, A. and Zhang, Y. (2017). Stock market mean reversion and portfolio choice over the life cycle. *Journal of Financial and Quantitative Analysis*, 52(3):1183–1209.

Michaelides, A. and Zhang, Y. (2022). Life-cycle portfolio choice with imperfect predictors. SSRN Working Paper.

Munk, C. and Sørensen, C. (2010). Dynamic asset allocation with stochastic income and interest rates. *Journal of Financial Economics*, 96(3):433–462.

Pflueger, C. E. and Viceira, L. M. (2011). Inflation-indexed bonds and the expectations hypothesis. *Annual Review of Financial Economics*, 3:139–158.

Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313.

Polk, C. and Vuolteenaho, T. (2026). Purifying the equity premium. SSRN Working Paper No. 5765362.

Rogoff, K. S., Rossi, B., and Schmelzing, P. (2024). Long-run trends in long-maturity real rates, 1311–2022. *American Economic Review*, 114(8):2271–2307.

Sangvinatsos, A. and Wachter, J. A. (2005). Does the failure of the expectations hypothesis matter for long-term investors? *Journal of Finance*, 60(1):179–230.

Shiller, R. J. (1989). *Market Volatility*. MIT Press, Cambridge, MA.

Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics*, 54(3):375–421.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717.

van Binsbergen, J. H. (2026). Duration-based stock valuation: Reassessing stock market performance and volatility. *Journal of Financial Economics*. NBER Working Paper No. 27367.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *Journal of Finance*, 56(2):433–470.

Wachter, J. A. (2002). Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets. *Journal of Financial and Quantitative Analysis*, 37(1):63–91.

A Calibration

Table A.1: Economic Calibration

Parameter	Symbol	Value
Panel A. Preferences and life-cycle		
Risk aversion (CRRA)	γ	5.0
Discount factor	β	0.96
Bequest horizon	$\bar{b}$	10 years
Bequest luxury shift	δ	0.0
Start age	–	22
Retirement age	–	67
Terminal age	–	99
Panel B. Deterministic age-earnings profile $g_t = b_0 + b_1\, t + b_2\, t^2/10 + b_3\, t^3/100$		
Intercept	b_0	-6.142
Linear	b_1	$+0.3040$
Quadratic	b_2	-0.051
Cubic	b_3	$+0.002586$
Panel C. Persistent income shock $z_{t+1} = \rho z_t + \eta_{t+1},\ \eta \sim p_z\, \mathcal{N}(\mu_{\eta 1}, \sigma^2_{\eta 1}) + (1-p_z)\, \mathcal{N}(\mu_{\eta 2}, \sigma^2_{\eta 2})$		
Persistence	ρ	0.991
Mixture weight	p_z	0.176
Component 1 mean	$\mu_{\eta 1}$	-0.524
Component 1 std	$\sigma_{\eta 1}$	0.113
Component 2 mean (derived)	$\mu_{\eta 2}$	0.112
Component 2 std	$\sigma_{\eta 2}$	0.046
Panel D. Transitory income shock $\varepsilon \sim p_\varepsilon\, \mathcal{N}(\mu_{\varepsilon 1}, \sigma^2_{\varepsilon 1}) + (1-p_\varepsilon)\, \mathcal{N}(\mu_{\varepsilon 2}, \sigma^2_{\varepsilon 2})$		
Mixture weight	p_ε	0.044
Component 1 mean	$\mu_{\varepsilon 1}$	$+0.134$
Component 1 std	$\sigma_{\varepsilon 1}$	0.762
Component 2 mean (derived)	$\mu_{\varepsilon 2}$	-0.0062
Component 2 std	$\sigma_{\varepsilon 2}$	0.055
Panel E. Payroll tax and Social Security AIME cap		
Payroll-tax rate (OASI)	τ_{SS}	10.6%
Payroll-tax / AIME cap	$\bar{y}_{\text{cap}}$	$2.5 \cdot \bar{L}$

Notes: All income quantities are in model units where the Social Security Average Wage Index equals one ($\bar{L} = 1$). The Component 2 means $\mu_{\eta 2}$ and $\mu_{\varepsilon 2}$ are pinned by the zero-mean restriction $\mu_{\cdot 2} = -\frac{p}{1-p}\, \mu_{\cdot 1}$ and computed in `build_model()` rather than treated as free parameters. Sources: preferences and bequest specification follow Catherine (2025); the deterministic age profile, persistent and transitory mixture-normal innovations follow Guvenen et al. (2022); the payroll tax matches the OASI portion of FICA. Bequest luxury shift δ is set to zero by the canonical solver override; see `configs/_canonical.py`.

Table A.2: Federal Income Tax Brackets (TCJA 2019, Single Filer)

Lower bound	Upper bound	Marginal rate
0.00	0.18	10%
0.18	0.72	12%
0.72	1.54	22%
1.54	2.94	24%
2.94	3.73	32%
3.73	9.32	35%
9.32	–	37%

Notes: Bracket bounds are expressed in units of the Social Security Average Wage Index $\bar{L}$. The schedule is applied to taxable income, defined as gross labor income net of the payroll tax. Pension income in retirement is taxed under the same schedule but is exempt from the payroll tax. Source: 2019 TCJA single-filer schedule.

Table A.3: Social Security Primary Insurance Amount (PIA) Schedule

Bracket	AIME range	Replacement rate
1	AIME ≤ 0.21	90%
2	$0.21 <$ AIME ≤ 1.25	32%
3	AIME > 1.25	15%

Notes: Average Indexed Monthly Earnings (AIME) approximated as $\text{AIME}(z) = \min\big(\exp(z) \cdot \overline{\exp(g_t)},\ 2.5\big)$, with the average taken over working ages and the cap matching the payroll-tax / taxable maximum. The PIA piecewise schedule applies bend points $b_1 = 0.21$ and $b_2 = 1.25$ with replacement rates $r_1 = 0.90$, $r_2 = 0.32$, $r_3 = 0.15$. The gross PIA is taxed under the same progressive schedule as labor income (Table A.2) but is exempt from payroll tax. Source: Catherine (2025), eqs. (17)–(20).

B Restricted-Predictor VAR Estimates

Table B.4: VAR(1) Estimates for System 1 (state = y_1) by Term Premium Intensity (λ)

	Predictor variables					
	$\mathrm{cape_t}$	$\mathrm{spr_t}$	$\mathrm{y\_1_t}$	$\mathrm{xr_t}$	$\mathrm{xb_t}$	R^2
Panel: $\lambda = 0$						
$\mathrm{cape_{t+1}}$	–	–	–	–	–	–
$\mathrm{spr_{t+1}}$	–	–	–	–	–	–
$\mathrm{y\_1_{t+1}}$	–	–	0.797* [0.521, 0.776]	–	–	0.628
$\mathrm{xr_{t+1}}$	–	–	−1.193 [−2.204, 0.044]	–	–	0.032
$\mathrm{xb_{t+1}}$	–	–	–	–	–	–
max $\lvert\mathrm{eig}(\Phi)\rvert = 0.797$						
Residual correlations						
cape	–					
spr	–	–				
y_1	–	–	1.00			
xr	–	–	0.13	1.00		
xb	–	–	−0.91	−0.06	1.00	
Panel: $\lambda = 0.5$						
$\mathrm{cape_{t+1}}$	–	–	–	–	–	–
$\mathrm{spr_{t+1}}$	–	–	–	–	–	–
$\mathrm{y\_1_{t+1}}$	–	–	0.797* [0.521, 0.776]	–	–	0.628
$\mathrm{xr_{t+1}}$	–	–	−1.193 [−2.204, 0.044]	–	–	0.032
$\mathrm{xb_{t+1}}$	–	–	−0.022 [−0.341, 0.086]	–	–	0.000
max $\lvert\mathrm{eig}(\Phi)\rvert = 0.797$						
Residual correlations						
cape	–					
spr	–	–				
y_1	–	–	1.00			
xr	–	–	0.13	1.00		
xb	–	–	−0.71	0.01	1.00	
Panel: $\lambda = 1$						
$\mathrm{cape_{t+1}}$	–	–	–	–	–	–
$\mathrm{spr_{t+1}}$	–	–	–	–	–	–
$\mathrm{y\_1_{t+1}}$	–	–	0.797* [0.521, 0.776]	–	–	0.628
$\mathrm{xr_{t+1}}$	–	–	−1.193 [−2.204, 0.044]	–	–	0.032
$\mathrm{xb_{t+1}}$	–	–	−0.144 [−0.404, 0.214]	–	–	0.004
max $\lvert\mathrm{eig}(\Phi)\rvert = 0.797$						
Residual correlations						
cape	–					
spr	–	–				
y_1	–	–	1.00			
xr	–	–	0.13	1.00		
xb	–	–	−0.46	0.07	1.00	

Notes: Point estimates with 90% bootstrap percentile interval [p_{05}, p_{95}] in square brackets. Here cape is the CAPE-based log earnings yield, $-\log(\mathrm{CAPE})$. In each panel, spr and xb use that panel's λ-specific long yield. Stationary block bootstrap of Politis and Romano (1994), applied to the annual primitive panel, mean block length 8, paired across λ. $B = 1000$ draws under the headline (conditional AR(1)) specification. Estimation sample: 1920–2011 ($T = 92$). * marks coefficients whose 90% interval excludes zero on the same side as the point estimate. max $\lvert\mathrm{eig}(\Phi)\rvert$ is the largest absolute eigenvalue of Φ. Stationarity requires < 1. Dashed entries (–) are zero by restriction or because the corresponding state is excluded from this restricted predictor system. At $\lambda = 0$ the xb equation is restricted to zero under the maintained EH benchmark ($\mathbb{E}_t[\mathrm{xb}_{t+1}] = 0$ in the return–state system). For persistent state coefficients, bootstrap percentile intervals may lie below the full-sample estimate because block joins attenuate estimated persistence in finite samples.

Table B.5: VAR(1) Estimates for System 2 (state = spr, y_1) by Term Premium Intensity (λ)

	Predictor variables					
	cape_t	spr_t	y_1_t	xr_t	xb_t	R^2
Panel: $\lambda = 0$						
cape_{t+1}	–	–	–	–	–	–
spr_{t+1}	–	0.446 [−1.021, 0.273]	−0.120* [−1.188, −0.231]	–	–	0.446
y_1_{t+1}	–	0.655* [0.806, 2.041]	1.096* [1.192, 2.101]	–	–	0.645
xr_{t+1}	–	1.327 [−27.953, 34.178]	−0.587 [−22.181, 23.095]	–	–	0.034
xb_{t+1}	–	–	–	–	–	–
max $\lvert\text{eig}(\Phi)\rvert = 0.935$						
Residual correlations						
cape	–					
spr	–	1.00				
y_1	–	−0.97	1.00			
xr	–	−0.17	0.12	1.00		
xb	–	0.82	−0.93	−0.06	1.00	
Panel: $\lambda = 0.5$						
cape_{t+1}	–	–	–	–	–	–
spr_{t+1}	–	0.594* [0.149, 0.679]	−0.013 [−0.229, 0.087]	–	–	0.368
y_1_{t+1}	–	0.374 [−0.048, 0.694]	0.949* [0.562, 1.031]	–	–	0.645
xr_{t+1}	–	0.757 [−4.393, 5.347]	−0.885 [−3.814, 1.332]	–	–	0.034
xb_{t+1}	–	1.056 [−0.193, 2.274]	0.407 [−0.416, 1.194]	–	–	0.037
max $\lvert\text{eig}(\Phi)\rvert = 0.935$						
Residual correlations						
cape	–					
spr	–	1.00				
y_1	–	−0.92	1.00			
xr	–	−0.19	0.12	1.00		
xb	–	0.50	−0.79	0.00	1.00	
Panel: $\lambda = 1$						
cape_{t+1}	–	–	–	–	–	–
spr_{t+1}	–	0.652* [0.317, 0.702]	0.049 [−0.121, 0.160]	–	–	0.369
y_1_{t+1}	–	0.261 [−0.024, 0.469]	0.890* [0.556, 0.902]	–	–	0.645
xr_{t+1}	–	0.530 [−2.506, 3.148]	−1.004 [−2.685, 0.210]	–	–	0.034
xb_{t+1}	–	1.471* [0.377, 2.395]	0.380 [−0.188, 1.076]	–	–	0.125
max $\lvert\text{eig}(\Phi)\rvert = 0.935$						
Residual correlations						
cape	–					
spr	–	1.00				
y_1	–	−0.88	1.00			
xr	–	−0.20	0.12	1.00		
xb	–	0.13	−0.58	0.06	1.00	

Notes: Point estimates with 90% bootstrap percentile interval [p_{05}, p_{95}] in square brackets. Here cape is the CAPE-based log earnings yield, $-\log(\text{CAPE})$. In each panel, spr and xb use that panel's λ-specific long yield. Stationary block bootstrap of Politis and Romano (1994), applied to the annual primitive panel, mean block length 8, paired across λ. $B = 1000$ draws under the headline (conditional AR(1)) specification. Estimation sample: 1920–2011 ($T = 92$). * marks coefficients whose 90% interval excludes zero on the same side as the point estimate. max $\lvert\text{eig}(\Phi)\rvert$ is the largest absolute eigenvalue of Φ. Stationarity requires < 1. Dashed entries (–) are zero by restriction or because the corresponding state is excluded from this restricted predictor system. At $\lambda = 0$ the xb equation is restricted to zero under the maintained EH benchmark ($\mathbb{E}_t[\text{xb}_{t+1}] = 0$ in the return–state system). For persistent state coefficients, bootstrap percentile intervals may lie below the full-sample estimate because block joins attenuate estimated persistence in finite samples.

1 Numerical convergence diagnostics

This appendix verifies that the discretized policy converges as the state grid is refined. The model is re-solved on five state grids of increasing resolution, $n^3 \in \{7^3,\ 8^3,\ 9^3,\ 10^3,\ 11^3\}$, with the persistent-income quadrature matched to the grid ($n_z \in \{16, 17, 18, 19, 20\}$). A sixth solve at 12^3 with $n_z = 20$ serves as the reference. All preferences, returns, demographics and VAR parameters are held at the calibration of the main text ($\gamma = 3,\ \lambda = 0$).

Table 1: Solver completion across grids

Grid	Sizes	n_z	Newton cells	Wall (min)	NaN	Inf	Newton convergence
g_7	7^3	16	42,291,900	91.4	0	0	99.94%
g_8	8^3	17	67,072,000	143.2	0	0	99.94%
g_9	9^3	18	101,112,300	216.6	0	0	99.94%
g_{10}	10^3	19	146,400,000	374.7	0	0	99.94%
g_{11}	11^3	20	205,107,100	439.7	0	0	99.94%
ref	12^3	20	266,284,800	572.3	0	0	99.94%

Notes: "Newton cells" is the total number of (age, state, persistent income state, wealth) cells at which the 2D Newton inner step is invoked. "Newton convergence" is the share of cells in which the inner Newton step reached the FOC tolerance within the iteration cap; for the residual 0.06% the EGM constrained-corner fallback is applied, and all such cells produced finite policies. The convergence share is constant across grids, indicating that the structural pocket of non-convergent cells lies in a fixed region of the state space (the kink near the borrowing-constrained corner) rather than reflecting an instability that worsens with refinement.

Table 2: Magnitude of the consumption-share policy change between consecutive refinements

Refinement step	Mean $\lvert\Delta(c/W)\rvert$	p99 $\lvert\Delta(c/W)\rvert$
$g_7 \to g_8$	0.0151	0.0614
$g_8 \to g_9$	0.0111	0.0495
$g_9 \to g_{10}$	0.0131	0.0551
$g_{10} \to g_{11}$	0.0103	0.0462
$g_{11} \to$ ref	**0.0003**	**0.0007**

Notes: Absolute cell-by-cell change in the consumption-share policy c/W between consecutive refinements, computed on a common state axis (median persistent income state, ages 25–95, wealth in $[0.1, 50]$ AWI). The first four refinement steps move the policy by roughly the same amount at every reported statistic. The final step ($g_{11} \to$ ref) collapses by a factor of ~ 30, indicating that the policy has entered the asymptotic regime in which further grid refinement no longer moves the solution.

Summary. Table 1 establishes that every grid resolution from g_7 through the 12^3 reference produced a complete, finite policy with the same 99.94% Newton convergence rate, with no NaN or Inf values anywhere in the solved policy. Table 2 and Figures 1–2 show that the policy itself stops moving as the grid is refined: the change between consecutive refinements is constant in order of magnitude through the first four steps and collapses by a factor of thirty at the final step. Taken together, these diagnostics confirm that the reference solution is converged in the sense of grid refinement.

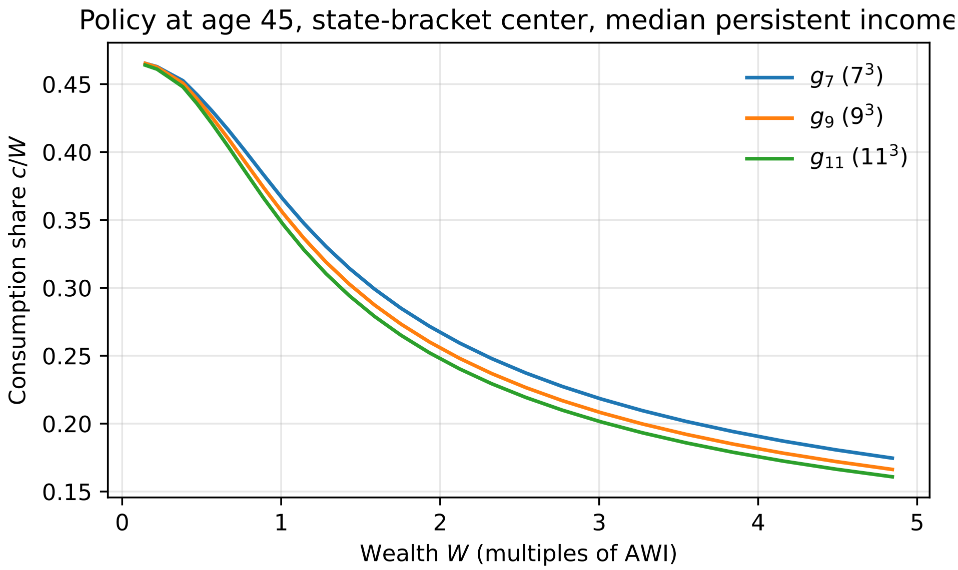

Figure 1: Consumption-share policy at age 45, evaluated at the state-bracket center and median persistent income state, for the odd-grid refinement sequence. The change in the policy between g_7 and g_9 is visibly larger than the change between g_9 and g_{11}, the pattern expected of a refinement sequence approaching its limit.

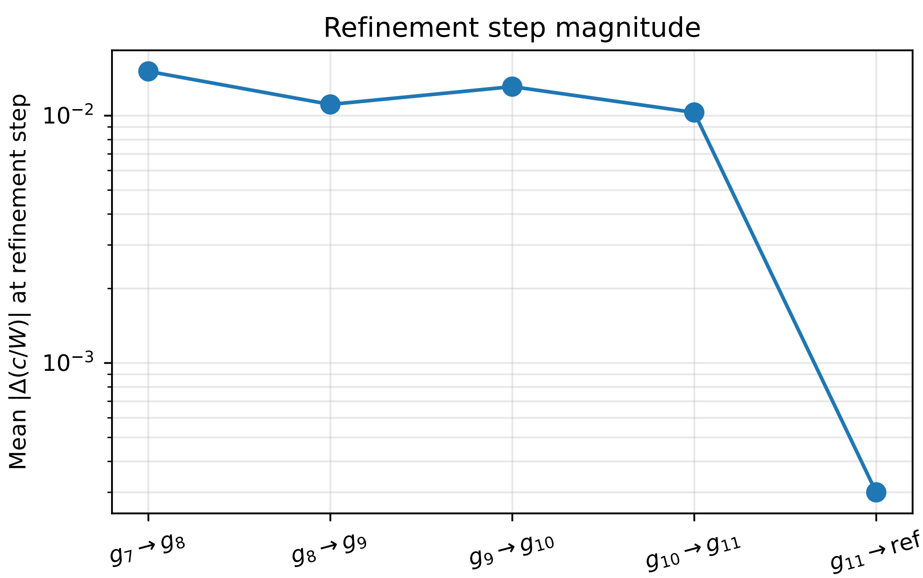

Figure 2: Mean absolute change in the consumption-share policy between consecutive refinements, log-scale. The first four refinements move the policy by a comparable amount; the final refinement step ($g_{11} \to$ ref) is approximately one order of magnitude smaller, the signature of the asymptotic regime having been reached.