

7519538-stock-bond-correlation-regimes-and-portfolio-choice-over-the-life-cycle

SUMMARY: This paper studies how the time-varying stock-bond correlation impacts optimal portfolio choices over the life-cycle. The authors embed a two-state Markov-switching process for the stock-bond correlation into an incomplete-markets life-cycle model, where households invest in cash, stocks, and long-term government bonds. The model shows that households infer the latent correlation regime from observed returns and adjust their portfolios based on their posterior beliefs. The paper finds that these belief-contingent portfolio differences are significant, especially near and during retirement, when financial wealth becomes a larger share of total lifetime resources. Despite the sizable allocation effects, the welfare gains from conditioning portfolios on covariance predictability are modest, with only 0.4 basis points of lifetime consumption improvement on average. However, households experiencing sustained periods of high posterior probability of the negative-correlation regime can see welfare gains of approximately 2.9 basis points.

METHOD: The study uses a two-state Markov-switching process for the stock-bond correlation embedded in an incomplete-markets life-cycle model. Households are assumed to invest in cash, stocks, and long-term government bonds, and their beliefs about the correlation regime are updated using Bayes' rule based on observed returns. The correlation process is estimated using a restricted Markov-switching maximum-likelihood specification that allows only the stock-bond correlation to vary across regimes.

KEY FINDINGS:

- Household portfolios are significantly affected by the inferred correlation regime.
- The magnitude of these effects varies over the life-cycle, being small for young households and larger near and during retirement.
- The welfare gains from conditioning portfolios on covariance predictability are modest, with only 0.4 basis points of lifetime consumption improvement on average.
- Households that experience sustained periods of high posterior probability of the negative-correlation regime can see substantial welfare gains of approximately 2.9 basis points.

IMPLICATION FOR TRADING: The findings suggest that traders and investors should consider the time-varying stock-bond correlation when making portfolio choices over the life-cycle. While the overall impact on portfolio allocation is significant, the welfare gains are relatively small. However, for households that experience prolonged periods of high posterior probability of the negative-correlation regime, the impact on portfolio allocation and welfare can be substantial. This implies that traders and investors should be aware of the changing correlation dynamics and adjust their portfolios accordingly, particularly near and during retirement, where financial wealth becomes a larger share of total lifetime resources.

Stock-Bond Correlation Regimes and Portfolio Choice over the Life-Cycle

September 2026

Abstract

The stock-bond correlation varies substantially over time, altering the diversification benefits of long-term government bonds. We embed a two-state Markov-switching process for the stock-bond correlation into an incomplete-markets life-cycle model in which households invest in cash, stocks, and long-term government bonds. Households infer the latent correlation regime from observed returns and condition their portfolios on posterior beliefs. The model generates sizable differences in optimal portfolios across beliefs, with persistence in the estimated correlation process playing a central role in their magnitude. These effects are concentrated near and during retirement, as financial wealth becomes a larger share of household resources. Despite these sizable allocation effects, conditioning portfolios on covariance predictability raises ex ante welfare by only 0.4 basis points of lifetime consumption on average. However, welfare gains are substantially larger for households that experience sustained periods of high posterior probability of the negative-correlation regime, reaching approximately 2.9 basis points for households with more than twenty years of such exposure.

Keywords: Life-Cycle Portfolio Choice; Stock-Bond Correlation; Bayesian Learning; Incomplete Markets

JEL Codes: G11, D91, G12, E21

1 Introduction

A common assumption in quantitative models of portfolio choice over the life-cycle is that the variance-covariance matrix of asset returns is constant. In the data, however, the correlation between major asset classes, especially stocks and long-term government bonds, varies substantially over time. As shown in Figure 1, decade-by-decade correlations between real annual returns on the S&P 500 and 10-year U.S. Treasury bonds range from strongly negative to strongly positive over the 1930–2025 sample period. Because the diversification benefits of long-term government bonds depend directly on their correlation with stocks, such shifts in correlation may matter for household portfolio choice.




Figure 1: Each bar reports the correlation between annual real returns on the S&P 500 and 10-year U.S. Treasury bonds within the indicated decade; the final bar reports the correlation for the 2020s to date. Source: Authors’ calculations based on [Damodaran \(2026\)](#), *Historical Returns on Stocks, Bonds and Bills: 1928–2025*; the 1928 and 1929 observations are excluded.

Motivated by Figure 1, this paper studies whether empirically plausible time-variation in the stock-bond correlation materially alters optimal household portfolio choice over the life-cycle. To proceed, we embed a two-state Markov-switching process for the stock-bond

correlation into an incomplete-markets life-cycle model in which households have access to a safe short-term asset with constant real return (hereafter referred to as “cash”), stocks, and long-term government bonds. The stochastic correlation process is estimated using a restricted Markov-switching maximum-likelihood (MLE) specification that holds return means and marginal volatilities fixed across regimes and allows only the stock-bond correlation to vary. This restriction isolates the correlation channel from broader changes in the return environment, allowing us to ask whether time-variation in stock-bond comovement is quantitatively important for portfolio choice over the life-cycle. Because the correlation regime is latent and thus not directly observable, households in the model use realized stock and bond returns to update their beliefs about the prevailing regime according to Bayes’ rule.

The calibrated model generates sizable belief-contingent portfolio differences. When the negative-correlation regime is more likely, long-term government bonds provide a stronger hedge against low stock-return states because their returns tend to be high when stock returns are low, leading households in the model to shift toward bonds and away from stocks and cash. For example, age 65 households with a high posterior probability of the negative-correlation regime allocate roughly 10 percentage points more of their financial wealth to long-term bonds, 6 percentage points less to stocks, and 4 percentage points less to cash than households with a high posterior probability of the positive-correlation regime.

The magnitudes of these portfolio responses vary meaningfully over the life-cycle. They tend to be small when households are young and financial wealth represents a relatively small share of total lifetime resources, and become substantially larger near and during retirement as financial wealth becomes a much larger share of total lifetime resources. In particular, among households in the model with positive savings, financial wealth rises from approximately 3 percent of total lifetime resources at age 35 to about 40 percent by age 65. Over the same age range, the difference in the long-term bond share between households with high negative-regime beliefs and high positive-regime beliefs rises from around 1 percentage point to approximately 10 percentage points. Thus, the life-cycle patterns of income and

wealth in the model determine when changes in beliefs about the prevailing stock-bond correlation regime become quantitatively important for household portfolio choice.

To better understand the source of these belief-contingent portfolio differences, we conduct a series of counterfactual exercises that separately vary regime persistence and long-run regime prevalence. Holding the stationary distribution across regimes fixed, portfolio responses remain quantitatively meaningful when regime persistence is reduced to the lower boundary of an empirically plausible region, but disappear when we consider an i.i.d. regime process since, in this case, current beliefs contain no information about future diversification opportunities. By contrast, equalizing the long-run prevalence of the two regimes while holding regime persistence fixed has only a modest effect on conditional allocations. These results indicate that regime persistence is the main feature of the transition process governing the magnitude of belief-contingent portfolio responses because greater persistence makes current beliefs more informative about future diversification opportunities.

Finally, we evaluate the welfare value of conditioning household portfolios on information about the correlation process by comparing the Bayesian belief-contingent policy rules with rules derived under an i.i.d. benchmark that preserves the stationary mixture of regimes but eliminates their predictive power. Both sets of policies are evaluated in the estimated Markov-switching environment. We find that the average welfare gain from exploiting covariance predictability is modest, at approximately 0.4 basis points of lifetime consumption, despite the sizable portfolio responses. However, the gains are substantially larger for households that experience sustained periods of high posterior probability of the negative-correlation regime, reaching approximately 2.9 basis points for households with more than twenty years of such exposure.

Our paper relates most closely to the large and growing literature that develops and analyzes quantitative models of portfolio choice over the life-cycle (e.g., [Cocco et al. \(2005\)](#), [Gomes and Michaelides \(2005\)](#), and [Benzoni et al. \(2007\)](#)). Much of this literature restricts the household portfolio to stocks and a riskless asset. A notable exception is [Tischbirek](#)

(2019), whose model allows households to allocate financial wealth across stocks, long-term government bonds, and a short-term riskless asset. While Tischbirek (2019) studies participation and portfolio composition over the life-cycle with an emphasis on the distinct roles played by long-term bonds and liquid safe assets as households accumulate and decumulate wealth, our paper instead focuses on persistent variation in the correlation between stocks and long-term government bonds. In particular, we examine how changes in the diversification value of long-term government bonds interact with the changing composition of household resources over the life-cycle when the prevailing correlation regime must be inferred from observed returns.

Our paper also connects to the literature on dynamic portfolio choice under time-varying investment opportunities (e.g., Kandel and Stambaugh (1996), Campbell and Viceira (1999), Viceira (2001), Campbell and Viceira (2002), Wachter (2002), and Liu (2011)). This literature shows that persistent variation in the conditional return distribution can materially affect portfolio demand, but much of it studies dynamic asset allocation outside a full quantitative life-cycle setting with labor income risk, retirement, mortality, and bequests. For example, Ang and Bekaert (2002) and Guidolin and Timmermann (2007) find large portfolio effects from regime switching in dynamic asset-allocation models with exogenous asset returns. We complement this literature by embedding persistent covariance variation and Bayesian learning about the prevailing correlation regime in a quantitative life-cycle model that allows us to study how the resulting portfolio responses vary over the life-cycle and across education groups and to assess their welfare implications.

Finally, our paper is related to the broader empirical literature documenting substantial time-variation in the stock-bond correlation and, therefore, in the hedging value of long-term government bonds (e.g., Ilmanen (2003), Connolly et al. (2005), Baele et al. (2010), Campbell et al. (2020), and Molenaar et al. (2024)). Relative to this literature, our contribution is to map this time-variation into a quantitative life-cycle model and ask whether our estimated stock-bond regime-switching process is persistent enough for beliefs about the

current regime to be predictive of future diversification opportunities and, thus, materially alter household portfolio choice.

The remainder of the paper proceeds as follows. Section 2 presents the model, including the household's Bayesian learning problem. Section 3 describes how we calibrate the model and estimate the stock-bond correlation regime-switching process. Section 4 presents our quantitative results. Finally, Section 5 concludes.

2 Model

Our model economy is populated by overlapping generations of households linked across generations through preferences, intergenerational transfers of wealth, education, and labor market ability. Households face labor income and lifespan risk and can allocate their wealth across stocks, long-term government bonds, and cash. The correlation between stock and long-term government bond returns evolves according to a latent two-state Markov-switching process. Households infer the current correlation regime from the realized history of stock and bond returns.

2.1 Preferences and Demographics

Households have CRRA preferences over consumption with coefficient of relative risk aversion γ and idiosyncratic subjective discount factor $\beta_i \in \{\beta_L, \beta_M, \beta_H\}$. Households enter the model at age 25 with exogenously given education level $j \in \{HD, HG, CG\}$, corresponding to high school dropout (HD), high school graduate (HG), and college graduate (CG), retire at age 65, and face survival risk thereafter up to a maximum age of 100, where $\phi_{t,t+1}$ denotes the probability of surviving from age t to age $t+1$ and we set $\phi_{100,101}$ equal to zero. A warm-glow bequest motive operates during retirement.

2.2 Income Process

Prior to retirement, household i with education level j at age t receives income endowment

$$y_{i,j,t} = \exp(f_{j,t} + \psi_{i,t} + e_{i,t}), \quad (1)$$

where $f_{j,t}$ is a deterministic education-specific age-earnings profile, $e_{i,t} \sim N(0, \sigma_{e,j}^2)$ is an i.i.d. transitory shock, and $\psi_{i,t}$ is a persistent component of income that evolves according to the random walk process

$$\psi_{i,t+1} = \psi_{i,t} + \varepsilon_{i,t+1}, \quad (2)$$

where $\varepsilon_{i,t+1} \sim N(0, \sigma_{\varepsilon,j}^2)$ is an i.i.d. persistent shock.¹

During retirement, households receive a non-stochastic income endowment equal to an education-specific replacement rate ξ_j of the persistent component of their income endowment at retirement:

$$y_{i,j,t} = \xi_j \exp(f_{j,65} + \psi_{i,65}). \quad (3)$$

2.3 Intergenerational Links

When households enter the model at age 25, they receive an immediate share Υ of their parent's bequest, inherit their parent's subjective discount factor $\beta_i \in \{\beta_L, \beta_M, \beta_H\}$ and persistent component of income $\psi_{i,65}$, and draw their education level j from the conditional distribution $G(\text{Education}|\text{Parent's Education})$. The remaining share of their parent's bequest is delayed. Beginning at age 45, households receive their delayed bequest with probability p each period. Households receive their delayed bequest at age 65 if they have not yet received it before then. To reduce our computational burden, we assume that delayed bequests are unanticipated by their recipients, both in terms of timing and size.² When a delayed bequest

¹For computational reasons, we assume that both transitory and persistent shocks to a household's income endowment are bounded and, in particular, that $e_{i,t} \in [\underline{e}_j, \bar{e}_j]$ and $\psi_{i,t} \in [\underline{\psi}_j, \bar{\psi}_j]$.

²De Nardi (2004) develops a model with intergenerational transfers of wealth in which bequest timing is stochastic and children infer the size of the bequest they are likely to receive by observing their parent's labor

is realized, it is added to the household's beginning-of-period wealth.

2.4 Assets and Stock-Bond Correlation Regimes

Households can allocate their savings across three assets: (i) a cash-like asset with constant real return r_c , (ii) stocks with real return $r_{s,t+1}$, and (iii) long-term government bonds with real return $r_{b,t+1}$. Following [Fagereng et al. \(2017\)](#), to limit stock market participation on the extensive margin, we impose a per-period fixed cost of holding stocks ζ .³

Let $z_t \in \{1, 2\}$ denote the latent stock-bond correlation regime at date t , where State 1 is the negative-correlation regime and State 2 is the positive-correlation regime. The stock-bond correlation regime evolves according to a two-state Markov process with transition matrix

$$\Pi = \begin{bmatrix} \pi_{1,1} & 1 - \pi_{1,1} \\ 1 - \pi_{2,2} & \pi_{2,2} \end{bmatrix}, \quad (4)$$

where $\pi_{n,m} = \Pr(z_{t+1} = m \mid z_t = n)$.

Conditional on next period's regime $z_{t+1} = k$, stock and bond returns are jointly normally distributed:

$$\begin{bmatrix} r_{s,t+1} \\ r_{b,t+1} \end{bmatrix} \Big| z_{t+1} = k \sim N \left(\begin{bmatrix} \mu_s \\ \mu_b \end{bmatrix}, \begin{bmatrix} \sigma_s^2 & \rho_k \sigma_s \sigma_b \\ \rho_k \sigma_s \sigma_b & \sigma_b^2 \end{bmatrix} \right), \quad (5)$$

where μ_s , μ_b , σ_s , and σ_b are constant across regimes and only the stock-bond correlation ρ_k depends on the realized regime. This specification isolates shifts in the stock-bond correlation from changes in average returns and marginal volatilities.

The correlation regime is not directly observed by households. Instead, households observe the history of realized stock and bond returns and form beliefs about the prevailing

productivity at entry into the model. A similar approach is taken by [De Nardi and Yang \(2014\)](#). However, this comes at the cost of increasing the model's state space and, therefore, the required computational time.

³This reduced-form parameter is intended to capture a range of barriers to equity market participation, including information acquisition costs, time and attention costs, account and advisory fees, and other institutional or behavioral frictions that are not modeled explicitly.

regime. Let

$$q_t \equiv \Pr(z_t = 1 \mid \mathcal{R}_t) \quad (6)$$

denote the posterior probability at date t that the economy is in the negative-correlation regime, where $\mathcal{R}_t$ denotes the history of stock and bond returns observed through date t . Given q_t , the predictive probability that next period's regime is State 1 is

$$\tilde{q}_{t+1} = q_t \pi_{1,1} + (1 - q_t)(1 - \pi_{2,2}). \quad (7)$$

After observing stock and bond returns at date $t+1$, households update their beliefs according to Bayes' rule:

$$q_{t+1} = \frac{\tilde{q}_{t+1} f_1(r_{s,t+1}, r_{b,t+1})}{\tilde{q}_{t+1} f_1(r_{s,t+1}, r_{b,t+1}) + (1 - \tilde{q}_{t+1}) f_2(r_{s,t+1}, r_{b,t+1})}, \quad (8)$$

where $f_k(\cdot)$ is the regime- k joint normal density implied by equation (5).⁴

Let $\Omega_t \equiv (\omega_{c,t}, \omega_{s,t}, \omega_{b,t})$ denote the portfolio-share vector chosen by households, which must satisfy

$$\omega_{c,t} + \omega_{s,t} + \omega_{b,t} = 1, \quad (9)$$

$$\omega_{c,t} \geq 0, \quad (10)$$

$$\omega_{s,t} \geq 0, \quad (11)$$

and

$$\omega_{b,t} \geq 0, \quad (12)$$

thus ruling out leverage and short-selling.

Let $\mathcal{A}(x_t)$ denote the feasible set for the portfolio share vector Ω_t , where x_t denotes household savings at date t . For households with positive savings (i.e., $x_t > 0$), $\mathcal{A}(x_t)$ is

⁴We assume households know the parameters governing the return process and form beliefs only about the current latent correlation regime.

the nonnegative simplex defined by equations (9)–(12). Households that choose not to save (i.e., $x_t \leq 0$) may elect to borrow in unsecured credit markets up to the natural borrowing constraint $\underline{x}_{j,t}$. Households that elect to borrow are restricted from investing in stocks, long-term bonds, or cash. Thus, $\mathcal{A}(x_t) = \{(0, 0, 0)\}$ if $x_t \leq 0$.

At date t , households choose the portfolio-share vector Ω_t conditional on q_t . At date $t+1$, the latent regime z_{t+1} is realized according to the transition process in equation (4), stock and bond returns are realized conditional on z_{t+1} according to equation (5), and households update their posterior belief to q_{t+1} according to equations (7) and (8).

For $x_t > 0$ and portfolio-share vector $\Omega_t \in \mathcal{A}(x_t)$ selected at date t , wealth at date $t+1$ is given by

$$a_{t+1} = x_t (1 + \omega_{c,t}r_c + \omega_{s,t}r_{s,t+1} + \omega_{b,t}r_{b,t+1}) - \zeta \mathbb{I}_{\omega_{s,t} > 0}, \quad (13)$$

where $\mathbb{I}_{\omega_{s,t} > 0}$ is an indicator function that is equal to one if $\omega_{s,t}$ is positive, and zero otherwise.

For $x_t \leq 0$, wealth at date $t+1$ is given by

$$a_{t+1} = x_t (1 + r_d), \quad (14)$$

where r_d is the real interest rate on unsecured debt.

2.5 Household Problem

The recursive problem of a working household at age $t \in \{25, 26, \dots, 64\}$ is⁵

$$V_{i,j,t}(a_t, \psi_t, e_t; q_t) = \max_{x_t \geq \underline{x}_{j,t}, \Omega_t \in \mathcal{A}(x_t)} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} + \beta_i E_t \left[V_{i,j,t+1}(a_{t+1}, \psi_{t+1}, e_{t+1}; q_{t+1}) \right] \right\}, \quad (15)$$

subject to equations (1), (2), (4) – (14), and the budget constraint

$$c_t + x_t = y_t + a_t, \quad (16)$$

⁵To simplify notation, we suppress subscripts where no confusion arises.

with expectations taken over the persistent and transitory labor income shocks and the joint predictive distribution of next period's latent correlation regime and asset returns.

During retirement, the recursive problem of a household at age $t \in \{65, 66, \dots, 100\}$ is

$$V_{i,j,t}(a_t, \psi_{65}; q_t) = \max_{x_t \geq \underline{x}_{j,t}, \Omega_t \in \mathcal{A}(x_t)} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} + \beta_i E_t \left[\phi_{t,t+1} V_{i,j,t+1}(a_{t+1}, \psi_{65}; q_{t+1}) + \dots \right. \right. \\ \left. \left. (1 - \phi_{t,t+1}) \nu(a_{t+1}) \right] \right\}, \quad (17)$$

subject to equations (3) – (14) and (16), with expectations taken over the joint predictive distribution of next period's latent correlation regime and asset returns. The bequest utility function is given by

$$\nu(a_{t+1}) = \left(\frac{\chi_1}{1 - \chi_1} \right)^\gamma \frac{[\chi_1(1 - \chi_1)^{-1} \chi_2 + a_{t+1}]^{1-\gamma}}{1 - \gamma} \quad (18)$$

as in Lockwood (2018), where $\chi_1 \in [0, 1)$ governs the strength of the bequest motive and $\chi_2 \geq 0$ determines the extent to which bequests are a luxury good.

3 Calibration

We divide the model parameters into two groups: those disciplined outside the model and those chosen jointly within the model to target moments in the data. We first discuss the parameters calibrated outside the model, including estimation of the stock-bond correlation process, before turning to the parameters jointly calibrated within the model.

3.1 Parameters Calibrated Outside the Model

Table 1 summarizes the parameters calibrated outside the model.

Table 1: Parameters calibrated outside the model

Parameter	Value	Source
<i>Preferences and Demographics</i>		
Coefficient of relative risk aversion (γ)	5	n/a
Subjective discount factor ($\beta_L, \beta_M, \beta_H$)	0.90, 0.94, 0.98	n/a
Survival probabilities ($\{\phi_{t,t+1}\}_{t=65}^{99}$)	U.S. Period Life Table (2019)	Arias and Xu (2022)
Delayed bequest probability (p)	0.05	see text
<i>Income Process</i>		
Age-Earnings Profiles		
High School Dropout, $f_{HD,t}$	$-2.14 + 0.168t - 3.5t^2/10^3 + 2.3t^3/10^5$	Cocco et al. (2005)
High School Graduate, $f_{HG,t}$	$-2.17 + 0.168t - 3.2t^2/10^3 + 2.0t^3/10^5$	Cocco et al. (2005)
College Graduate, $f_{CG,t}$	$-4.32 + 0.319t - 5.8t^2/10^3 + 3.3t^3/10^5$	Cocco et al. (2005)
Income-Shock Variances		
High School Dropout, $(\sigma_{\varepsilon,HD}^2, \sigma_{e,HD}^2)$	0.0105, 0.1056	Cocco et al. (2005)
High School Graduate, $(\sigma_{\varepsilon,HG}^2, \sigma_{e,HG}^2)$	0.0106, 0.0738	Cocco et al. (2005)
College Graduate, $(\sigma_{\varepsilon,CG}^2, \sigma_{e,CG}^2)$	0.0169, 0.0584	Cocco et al. (2005)
Retirement-Income Replacement Rates		
High School Dropout, ξ_{HD}	0.89	Cocco et al. (2005)
High School Graduate, ξ_{HG}	0.68	Cocco et al. (2005)
College Graduate, ξ_{CG}	0.94	Cocco et al. (2005)
<i>Asset Returns and Correlation Regimes</i>		
Cash-like asset real return, r_c	0.004	Sample moment
Common stock mean and volatility, (μ_s, σ_s)	(0.086, 0.193)	Sample moments
Common bond mean and volatility, (μ_b, σ_b)	(0.018, 0.089)	Sample moments
Regime correlations, (ρ_1, ρ_2)	(-0.81, 0.25)	Restricted MLE
Transition probabilities, $(\pi_{1,1}, \pi_{2,2})$	(0.903, 0.981)	Restricted MLE

Notes: HD, HG, and CG denote high school dropout, high school graduate, and college graduate, respectively. State 1 is the negative-correlation regime and State 2 is the positive-correlation regime. In the restricted specification, common stock and bond means and marginal volatilities are fixed at their unconditional full-sample estimates.

Each model period represents one calendar year. We set the coefficient of relative risk aversion γ equal to 5, a standard value in the quantitative life-cycle portfolio choice literature, and allow the idiosyncratic subjective discount factors β_L , β_M , and β_H to take on values of 0.90, 0.94, and 0.98, respectively. Survival probabilities $\{\phi_{t,t+1}\}_{t=65}^{99}$ are taken from the 2019 U.S. period life table for the total population (Arias and Xu (2022)) with $\phi_{100,101}$ set equal to zero. The annual delayed bequest probability p is set equal to 0.05, and households receive their delayed bequest at age 65 if they have not yet received it before then.⁶ We calibrate the process governing the intergenerational transmission of educational

⁶This implies a truncated geometric distribution over the timing of delayed bequest receipt, with receipt probability $p(1-p)^{t-45}$ at each age $t \in \{45, \dots, 64\}$ and a terminal mass of $(1-p)^{20} \approx 0.358$ at age 65.

attainment, represented by the conditional distribution $G(\text{Education}|\text{Parent's Education})$, using data from the U.S. Department of Education's Longitudinal Study of 2002 (Table 2). We parameterize the labor-income process following Cocco et al. (2005), including the education-specific age-earnings profiles, variances of persistent and transitory income shocks, and retirement-income replacement rates.

Table 2: Educational Attainment Conditional on Parent's Education

Parent's Education	Child's Education		
	HD	HG	CG
HD	28%	58%	14%
HG	16%	57%	27%
CG	6%	38%	56%

Notes: HD, HG, and CG denote high school dropout, high school graduate, and college graduate, respectively. Source: U.S. Department of Education, NCES, Longitudinal Study of 2002, High School Sophomores.

We discipline asset return moments using Aswath Damodaran's *Historical Returns on Stocks, Bonds and Bills: 1928–2025* (Damodaran (2026)). The real annual return on the cash-like asset is set to $r_c = 0.004$, matching the average real return on 3-month U.S. Treasury bills. For risky assets, we use annual real returns on the S&P 500 and 10-year U.S. Treasury bonds. Over the full sample, the unconditional mean and standard deviation of real annual stock returns are $(\mu_{s,u}, \sigma_{s,u}) = (0.086, 0.193)$, while the corresponding moments for real annual bond returns are $(\mu_{b,u}, \sigma_{b,u}) = (0.018, 0.089)$. Over the full sample, the unconditional stock-bond correlation is $\rho_u = 0.09$.

Next, we turn to disciplining the Markov-switching process for time-variation in the stock-bond correlation. Our annual sample contains 98 return observations, which limits the precision with which a flexible latent-state process can be estimated. In particular, a specification with state-dependent means, variances, and correlations is weakly identified and sensitive to modeling choices since estimating separate means and volatilities by state leaves relatively little information for distinguishing persistent covariance variation from

transient return realizations. In exploratory estimates, such flexibility produces unstable transition probabilities and implausibly volatile regime-specific moments.⁷ We therefore adopt a restricted two-state Markov-switching specification as a disciplined reduced-form summary of persistent covariance variation, rather than treating the data as sufficient to identify a rich structural state process. In this specification, stock and bond means and marginal volatilities are common across states and only their correlation varies.⁸

We proceed in two steps. First, we set the common stock and bond means and marginal volatilities equal to their unconditional full-sample estimates, $(\mu_s, \sigma_s) = (0.086, 0.193)$ and $(\mu_b, \sigma_b) = (0.018, 0.089)$, respectively. Second, taking these moments as fixed, we maximize the two-state Gaussian Markov-switching likelihood over the regime-dependent stock-bond correlations, ρ_1 and ρ_2 , and the transition probabilities, $\pi_{1,1}$ and $\pi_{2,2}$.

Table 3 reports 90 percent confidence intervals for the return-process parameters. Confidence intervals for the unconditional means and marginal volatilities are computed using Newey–West standard errors. Uncertainty in the state-dependent correlations and transition probabilities is summarized using one-parameter profile-likelihood confidence intervals computed while holding the unconditional means and marginal volatilities fixed at their full-sample estimates.

The estimated state-dependent stock-bond correlations differ sharply across states, with $\rho_1 = -0.81$ in the negative-correlation regime and $\rho_2 = 0.25$ in the positive-correlation regime. The estimated transition probabilities imply substantial persistence in both regimes,

⁷In particular, we find that a fully unrestricted specification with state-dependent means, volatilities, and correlations is poorly behaved in the short annual sample and tends toward boundary solutions. Our restricted specification serves not only as a simplification, but also as a disciplining device that focuses the latent state on the stock-bond correlation channel rather than on unstable state-specific means and volatilities. We also explore a partially unrestricted specification that allows stock and bond means, the stock-bond correlation, and transition probabilities to vary across regimes while maintaining common stock and bond marginal volatilities. Importantly, the resulting point estimates for the correlations and transition probabilities are nearly identical to those from the baseline restricted MLE. Moreover, the smoothed regime probabilities imply very similar regime timing. When we re-solve and re-simulate our calibrated model under the partially unrestricted estimates, the direction of the belief-contingent portfolio response is preserved but its magnitude increases substantially. See Appendix C for details.

⁸Because the regime-specific means and marginal volatilities are common across states, differences in model-implied belief-contingent portfolios are not driven by differences in expected returns or marginal risk.

with $(\pi_{1,1}, \pi_{2,2}) = (0.903, 0.981)$, a key condition for current beliefs about the latent regime to contain information about future diversification opportunities relevant for long-term investors. However, the 90 percent profile-likelihood interval for the persistence of the negative-correlation regime spans $[0.70, 0.98]$, indicating substantial uncertainty about $\pi_{1,1}$. The persistence of the positive-correlation regime, on the other hand, is more precisely identified, as the 90 percent profile-likelihood interval for $\pi_{2,2}$ spans $[0.92, 1.00]$.

Table 3: Statistical Uncertainty for Estimated Return-Process Parameters

Parameter	Estimate	90 Percent Confidence Interval
Stock mean, μ_s	0.086	[0.058, 0.115]
Stock volatility, σ_s	0.193	[0.169, 0.217]
Bond mean, μ_b	0.018	[0.003, 0.034]
Bond volatility, σ_b	0.089	[0.076, 0.101]
Negative-regime correlation, ρ_1	-0.81	[-0.90, -0.56]
Positive-regime correlation, ρ_2	0.25	[0.07, 0.40]
Negative-regime persistence, $\pi_{1,1}$	0.903	[0.696, 0.984]
Positive-regime persistence, $\pi_{2,2}$	0.981	[0.916, 0.998]

Notes: State 1 is the negative-correlation regime and State 2 is the positive-correlation regime. The stock and bond means and marginal volatilities are unconditional full-sample estimates and are held fixed in the restricted Markov-switching likelihood. Conditional on these moments, the likelihood estimates the regime-dependent stock-bond correlations, ρ_1 and ρ_2 , and the transition probabilities, $\pi_{1,1}$ and $\pi_{2,2}$. Confidence intervals for the unconditional means and marginal volatilities are computed using Newey–West standard errors. Confidence intervals for ρ_1 , ρ_2 , $\pi_{1,1}$, and $\pi_{2,2}$ are one-parameter profile-likelihood confidence intervals computed while holding the unconditional stock and bond means and marginal volatilities fixed at their full-sample estimates.

Figure 2 plots the profile log-likelihood loss over $(\pi_{1,1}, \pi_{2,2})$ under the same common-moment restrictions as the baseline MLE. The likelihood surface places considerably more support on persistent covariance regimes than on low-persistence alternatives, although the joint 90 percent likelihood-ratio region remains broad, particularly along the dimension governing persistence of the negative-correlation regime.

At the point estimates, the implied average regime durations are 10.3 years for the negative-correlation state and 51.9 years for the positive-correlation state, with station-




Figure 2: Profile likelihood loss for the transition probabilities, $\ell(\hat{\theta}) - \ell_p(\pi_{1,1}, \pi_{2,2})$, under the restricted Markov-switching specification. The heatmap is evaluated over a finite grid of transition probabilities, while the red marker denotes the exact MLE obtained from continuous optimization. The contour line denotes the joint 90 percent likelihood-ratio confidence region for $(\pi_{1,1}, \pi_{2,2})$. At each grid point, the transition probabilities are fixed and the two regime-dependent correlations are re-optimized, while the common stock and bond means and marginal volatilities are held fixed at their unconditional full-sample estimates. State 1 is the negative-correlation regime and State 2 is the positive-correlation regime. The blue square denotes the low-persistence counterfactual used in Section 4.4, and the dashed line traces the path from this counterfactual to the restricted MLE along which the stationary distribution is held fixed.

ary probabilities of 0.166 and 0.834, respectively.⁹ Thus, the positive-correlation regime is both highly persistent conditional on entry and dominant in the long-run stationary distribution. These quantities should nevertheless be viewed as useful summaries of the fitted restricted process rather than precise estimates of regime duration or prevalence.¹⁰ Given the uncertainty in our estimates of regime persistence, Section 4.4 examines how the portfolio implications of our Bayesian-learning model change under alternative transition dynamics. Figure 2 marks the transition matrix used in Section 4.4, chosen just inside the low-persistence boundary of the joint 90 percent likelihood-ratio region along a path that holds the stationary distribution fixed.

Panel (a) of Figure 3 plots the filtered and smoothed probabilities that the economy is in the positive-correlation regime. The smoothed probability uses the full sample of returns to infer the regime prevailing at each date, while the filtered probability uses only returns observed through that date. The filtered series therefore corresponds to the real-time state inference problem faced by households in our quantitative model, conditional on the estimated return-process parameters.¹¹ Both probability series broadly track the low-frequency variation in stock-bond comovement documented in Figure 1. Filtered beliefs are relatively concentrated during most of the sample, although they become substantially more diffuse around regime transitions. Indeed, the maximum filtered state probability exceeds 0.80 in 79.6 percent of sample years and 0.90 in 67.3 percent. Panel (b) of Figure 3 maps these filtered beliefs into the corresponding one-period-ahead predictive stock-bond correlation.

⁹For a two-state Markov chain with transition matrix

$$\Pi = \begin{bmatrix} \pi_{1,1} & 1 - \pi_{1,1} \\ 1 - \pi_{2,2} & \pi_{2,2} \end{bmatrix},$$

the expected duration of state i is $1/(1 - \pi_{i,i})$. The stationary distribution is given by

$$\bar{\pi}_1 = \frac{1 - \pi_{2,2}}{2 - \pi_{1,1} - \pi_{2,2}}, \quad \bar{\pi}_2 = \frac{1 - \pi_{1,1}}{2 - \pi_{1,1} - \pi_{2,2}}.$$

¹⁰The partially unrestricted specification in Appendix C produces similar correlations, transition probabilities, and regime timing.

¹¹We hold the estimated return-process parameters fixed at their full-sample estimates rather than recursively re-estimating them as the sample expands.




Figure 3: Filtered and smoothed probabilities of the positive-correlation regime, and the corresponding filtered one-period-ahead predictive stock-bond correlation, under the restricted two-state Markov-switching specification. Panel (a) plots the smoothed probability $P(z_t = 2 \mid R_1, \dots, R_T)$ and the filtered probability $P(z_t = 2 \mid R_1, \dots, R_t) = 1 - q_t$ of the positive-correlation regime, where q_t is the posterior probability of the negative-correlation regime defined in Section 2. Panel (b) plots the filtered one-period-ahead predictive stock-bond correlation, computed as $P(z_{t+1} = 1 \mid R_1, \dots, R_t)\rho_1 + P(z_{t+1} = 2 \mid R_1, \dots, R_t)\rho_2$. The horizontal dotted lines indicate the full-information one-period-ahead predictive correlations associated with the negative- and positive-correlation regimes. Filtered probabilities use only returns observed through year t for state inference, while the parameters governing the return process are held fixed at their full-sample estimates.

During extended periods of regime stability, the predictive correlation approaches the full-information value associated with the prevailing regime, while around regime transitions it adjusts more gradually as beliefs are updated based on observed returns. Taken together, these patterns indicate that observed returns are generally informative about the latent correlation regime, while uncertainty rises around regime transitions when the current state is more difficult to infer.

3.2 Parameters Calibrated Within the Model

The parameters calibrated within the model include the distribution of households over subjective discount factors $\eta_k \equiv \Pr(\beta_i = \beta_k)$ for $k = L, M, H$, the interest rate on unsecured debt r_d , the per-period cost of holding stocks ζ , the fraction of bequests made available immediately to heirs Υ , and the bequest parameters χ_1 and χ_2 . We calibrate these parameters jointly to target corresponding moments in the data.

We set the probability that a household has the subjective discount factor β_M equal to 0.62 to target the ratio of median wealth to median income of 1.95.¹² We set the probability that a household has the subjective discount factor β_H equal to 0.16 to target the ratio of wealth held by households at the 75th percentile of the distribution to median income of 6.25.¹³ The probability that a household has the subjective discount factor β_L is set to 0.22 so that the probabilities sum to one. We set the interest rate on unsecured debt r_d equal to 32.5% to target the ratio of wealth held by households at the 10th percentile of the distribution to median income of approximately 0.¹⁴ We set the per-period fixed cost of holding stocks (relative to median income) ζ equal to 0.073 to target a stock market

¹²The median household wealth target of \$115,320 is reported by [Ashman and Neumuller \(2020\)](#) using observations pooled across the 1989–2016 waves of the Survey of Consumer Finances (SCF), with wealth expressed in 2016 dollars. We normalize this wealth measure by the 2016 U.S. median household income of \$59,039 reported by the U.S. Census Bureau using the Current Population Survey Annual Social and Economic Supplement (CPS ASEC) ([Semega et al., 2017](#)).

¹³Using observations pooled across the 1989–2016 waves of the SCF, [Ashman and Neumuller \(2020\)](#) report household net wealth at the 75th percentile of the distribution equal to \$368,943 in 2016 dollars.

¹⁴Using observations pooled across the 1989–2016 waves of the SCF, [Ashman and Neumuller \(2020\)](#) report household net wealth at the 10th percentile of the distribution equal to -\$281 in 2016 dollars.

participation rate of 58%.¹⁵ The fraction of bequests made available immediately to heirs Υ is set equal to 0.21 to target the ratio of wealth held by age 25 households to median income of 0.32.¹⁶ Finally, we set the bequest parameters χ_1 and χ_2 equal to 0.846 and 9.10, respectively, in order to target the ratio of wealth held by households at the 90th percentile of the distribution to median income of approximately 16.8 and the ratio of annual bequests to total household wealth of roughly 0.018.¹⁷

Table 4: Calibrated Parameters and Model Fit

Parameter	Value	
Distribution over subjective discount factors (η_L, η_M, η_H)	0.22, 0.62, 0.16	
Interest rate on household debt (r_d)	32.5%	
Fixed cost of holding stocks (ζ)	0.073	
Fraction of bequest inherited at age 25 (Υ)	0.21	
Bequest parameters (χ_1, χ_2)	0.846, 9.10	
Moment	Data	Model
10th percentile	0.00	0.00
50th percentile	1.95	1.95
75th percentile	6.25	6.75
90th percentile	16.8	16.7
Median household wealth at age 25	0.32	0.24
Stock market participation rate	0.58	0.58
Ratio of annual bequests to total household wealth	0.018	0.017

Notes: Wealth is normalized by median household income.

As shown in Table 4, the calibrated model matches most of the target moments reasonably well, although it understates median household wealth at age 25 and overstates household wealth at the 75th percentile of the distribution.¹⁸ The relatively high value for the fixed

¹⁵<https://www.sec.gov/data-research/statistics-data-visualizations/us-households-participation-capital-markets>, accessed April 28, 2026.

¹⁶Using observations pooled across the 1989–2016 waves of the SCF, Ashman and Neumuller (2020) report households headed by an age 25 individual had net wealth equal to \$18,864 in 2016 dollars.

¹⁷Using observations pooled across the 1989–2016 waves of the SCF, Ashman and Neumuller (2020) report household net wealth at the 90th percentile of the distribution equal to \$990,942 in 2016 dollars. Kotlikoff and Summers (1981) find that intergenerational transfers account for the vast majority of aggregate wealth in the U.S. In particular, the authors estimate the stock of transfer wealth in the U.S. was \$3.151 trillion in 1974. The authors then estimate that, in order to account for this vast sum of wealth, the annual flow of transfers would need to be approximately \$70 billion. Given a total net worth for the household sector of \$3.884 trillion in 1974, the implied ratio of annual intergenerational transfers of wealth to the stock of household wealth is 0.0180.

¹⁸The 32.5 percent interest rate on unsecured debt primarily affects the level and curvature of the wealth accumulation profile early in the working life, rather than portfolio composition conditional on positive savings. Because our portfolio-share results are computed conditional on households with positive savings,

cost of holding stock partly reflects the 8.2 percent equity premium in the model, which otherwise generates a strong incentive to participate in the stock market.¹⁹ In Appendix E we show that our main results are robust to setting ζ equal to 2% of median household income, a value substantially lower than the baseline calibration.

4 Results

We now assess the quantitative importance of time-variation in the stock-bond correlation for portfolio choice and welfare over the life-cycle. First, we consider fixed negative- and positive-correlation regimes. These benchmarks shut down regime-switching risk and learning about the latent correlation regime in order to measure the pure effect of changing the stock-bond correlation on portfolio allocations while holding expected returns and marginal volatilities fixed. Second, we introduce the estimated Markov-switching process under the assumption that households observe the prevailing correlation regime. This full-information benchmark shows whether the portfolio differences from the fixed-regime benchmarks remain when the correlation state is persistent rather than permanent. Third, we turn to our baseline model in which households condition their portfolio decisions on posterior beliefs about the latent correlation regime. Fourth, we perform counterfactual exercises that separately vary regime persistence and long-run regime prevalence to determine how each feature of the transition process affects belief-contingent portfolio choice over the life-cycle. Finally, we quantify the welfare value of allowing households to condition their portfolio decisions on information about the prevailing stock-bond correlation regime.

For the portfolio results in Sections 4.1–4.3, we simulate 10 million households facing labor income risk, education heterogeneity, mortality risk, and stochastic stock and bond returns.²⁰

the portfolio implications are not mechanically driven by this parameter since households with negative or near-zero wealth are effectively at a corner solution (i.e., no risky asset holdings) regardless of their beliefs about the correlation regime.

¹⁹Since the cash-like asset is the risk-free asset in our model, the equity premium is given by $\mu_s - r_c = 0.086 - 0.004 = 0.082$.

²⁰Appendix A provides additional computational details.

Unless otherwise noted, figures and tables report equally weighted average portfolio shares by age among households with positive savings. In Sections 4.1 and 4.2, these averages are conditional on the current stock-bond correlation regime. In Section 4.3, they are conditional on households' posterior beliefs about the latent regime.

4.1 Fixed-Regime Benchmarks

We begin by fixing the economy permanently in either the negative- or positive-correlation regime. In these full-information benchmarks, the correlation environment is known and constant over time. Because stock and bond means and marginal volatilities are held fixed across regimes, any differences in portfolio allocations are driven entirely by differences in the stock-bond correlation.

Figure 4 depicts the resulting life-cycle portfolio allocations for each fixed regime. In both cases stock shares rise over much of the working life while bond shares decline. These broad life-cycle profiles reflect the interaction of wealth accumulation, labor-income risk, and the fixed cost of stock-market participation. Early in the working life, financial portfolios are relatively small and the participation cost is large relative to financial wealth. As households accumulate financial wealth and remaining labor-income risk becomes smaller relative to their financial portfolios, stock shares rise and bond shares decline. Stock shares subsequently fall and allocations shift back toward bonds and cash as households progress through retirement and decumulate financial wealth.

The key comparison for our purposes is across the two correlation regimes. In the negative-correlation regime, long-term bonds provide meaningful diversification benefits, so households hold substantially more bonds near and during retirement, less stock, and very little cash. Although cash provides a (low) risk-free real return, it does not provide the same state-contingent hedge against equity risk as long-term government bonds. The lower cash share in the negative-correlation regime therefore partly reflects the stronger hedging value of long-term government bonds within the three-asset portfolio. In the positive-correlation




Figure 4: Life-cycle portfolio shares, equally weighted across all households with positive savings, when the economy is permanently fixed in the negative-correlation regime (solid black) or permanently fixed in the positive-correlation regime (dashed red).

regime, by contrast, the diversification benefits of long-term government bonds are smaller, so households substitute away from bonds and toward stock and cash, especially as they get close to and enter into retirement.

Table 5 shows that differences in portfolio holdings between the negative- and positive-correlation regimes are economically large. At ages 55, 65, and 75, households in the fixed negative-correlation regime hold roughly 4 to 7 percentage points less stock, 6 to 11 percentage points more bonds, and 1 to 4 percentage points less cash than households in the fixed positive-correlation regime. These results suggest that, through the lens of our quantitative model, differences in stock-bond correlation alone are capable of generating materially different life-cycle portfolios across fixed negative- and positive-correlation regimes.

Table 5: Stock, Bond, and Cash Shares: Fixed Correlation Environments

Scenario	Stock Share			Bond Share			Cash Share		
	55	65	75	55	65	75	55	65	75
Fixed Negative Correlation	0.74	0.71	0.60	0.26	0.28	0.38	0.01	0.01	0.02
Fixed Positive Correlation	0.78	0.78	0.66	0.20	0.17	0.29	0.02	0.05	0.06

Notes: These full-information benchmarks isolate the diversification effect of moving between permanently fixed negative- and positive-correlation environments. Portfolio shares are equal-weighted among households with positive savings.

4.2 Regime Switching Under Full Information

We next allow the correlation regime to evolve according to the estimated transition process and ask whether these portfolio differences across regimes remain quantitatively important when the current correlation regime is persistent rather than permanent.

Figure 5 shows the differences in portfolio shares between households in the negative- and positive-correlation regimes in this full-information regime-switching benchmark. The regime-conditional differences remain economically large and closely resemble those obtained in the fixed-regime experiment, particularly near and during retirement. Households continue to hold substantially more long-term government bonds and less stock and cash in the negative-correlation regime than in the positive-correlation regime. The differences are

relatively small early in the working life, increase substantially as households approach retirement, and remain large during retirement.




Figure 5: Differences in life-cycle portfolio shares between the negative- and positive-correlation regimes under the full-information Markov-switching model. Each line reports the negative-regime portfolio share minus the positive-regime portfolio share, in percentage points, equally weighted across households with positive savings.

Table 6 reports the corresponding regime-conditional stock, bond, and cash shares at selected ages. At ages 55, 65, and 75, households in the negative-correlation regime allocate approximately 6, 10, and 10 percentage points more of their financial portfolios to long-term government bonds, respectively, than households in the positive-correlation regime. These higher bond holdings are accompanied by lower allocations to stocks and cash. The gap

between the negative- and positive-correlation allocations is therefore large not only in the fixed-regime experiment, but also conditional on the current regime in this full-information regime-switching benchmark. The similarity between the fixed-regime and regime-switching results is consistent with the substantial persistence of the estimated correlation process. Although the current regime is not perfectly persistent, it is persistent enough to remain informative about the future stock-bond correlation and therefore about the future diversification benefits of long-term government bonds.

Table 6: Stock, Bond, and Cash Shares under Full-Information Regime Switching

Scenario	Stock Share			Bond Share			Cash Share		
	55	65	75	55	65	75	55	65	75
Negative Correlation State	0.74	0.71	0.60	0.26	0.28	0.39	0.01	0.00	0.02
Positive Correlation State	0.78	0.78	0.66	0.20	0.18	0.29	0.02	0.04	0.05

Notes: The comparison isolates portfolio differences between the negative- and positive-correlation states under the estimated persistent regime process. Portfolio shares are equal-weighted among households with positive savings.

We next ask whether these regime-contingent portfolio differences survive when households do not observe the current correlation regime directly and instead must infer it from the history of realized stock and bond returns.

4.3 Bayesian Learning and Portfolio Choice

In our baseline model, households use observed stock and bond returns to update their posterior probability of the negative-correlation regime according to Bayes' rule.

To characterize the resulting portfolio differences, we compare households with a high posterior probability of the negative-correlation regime, defined as $q_t \geq 0.8$, to households with a high posterior probability of the positive-correlation regime, defined as $q_t \leq 0.2$.²¹

These posterior beliefs are relatively concentrated in the simulated model. As shown in Panel

²¹These are simulated conditional averages and therefore reflect both the direct response of household portfolio policies to beliefs and any differences in the distribution of other household states across the two belief groups. In Section 4.4, we isolate the direct policy response by holding the distribution of non-belief household states fixed.

A of Table 7, households assign at least 80 percent probability to one of the two regimes in approximately 86 percent of simulated observations and at least 90 percent probability in approximately 72 percent of observations. Moreover, the regime associated with the larger posterior probability coincides with the underlying latent state approximately 92 percent of the time. Thus, realized stock and bond returns provide households with substantial, although imperfect, information about the prevailing correlation regime.²²

Table 7: Bayesian Learning and Portfolio Responses

Panel A: Posterior diagnostics				
Diagnostic				Value
High-negative posterior, $q_t \geq 0.8$				8.1%
High-positive posterior, $q_t \leq 0.2$				77.6%
Diffuse posterior, $q_t \in (0.2, 0.8)$				14.3%
Posterior confidence ≥ 0.80				85.7%
Posterior confidence ≥ 0.90				72.2%
Posterior confidence ≥ 0.95				56.5%
Latent-state classification accuracy				91.8%
Brier score				0.063
Panel B: Portfolio-share responses				
Age	Full Information Bond Gap	Bayesian-Learning Bond Gap	Bayesian-Learning Stock Gap	Bayesian-Learning Cash Gap
55	5.9	5.8	-4.4	-1.4
65	10.6	10.0	-6.5	-3.5
75	9.7	9.4	-6.3	-3.0

Notes: $q_t = \Pr(z_t = 1 \mid \mathcal{R}_t)$ is the posterior probability of the negative-correlation regime. The Bayesian-learning portfolio-share gaps in Panel B are the average shares among households with $q_t \geq 0.8$ minus the corresponding average shares among households with $q_t \leq 0.2$, in percentage points, restricting attention to households with positive savings. The full-information bond gap is the negative-state minus positive-state difference in average bond shares under the matched full-information Markov-switching benchmark, in which households observe the current latent regime z_t directly. The Bayesian solution uses 11 belief points and the same 100 regime-specific return nodes as the full-information benchmark.

The relatively high concentration of posterior beliefs should not be interpreted as a generic implication of learning about covariance regimes. In our estimated process, the two regime correlations are widely separated, at -0.81 and 0.25 , and both regimes are persistent. The resulting concentration of posterior beliefs therefore reflects the combination of strong regime

²²Appendix B reports numerical convergence of the belief-contingent portfolio responses with respect to the discretization of the posterior-belief state.

separation, persistent state dynamics, and the assumption that households know the parameters governing the return process. Learning would generally be less precise if the regimes were harder to distinguish, less persistent, or if households also faced uncertainty about the return-process parameters.




Figure 6: Differences in life-cycle portfolio shares between households with high posterior probability of the negative-correlation regime and households with high posterior probability of the positive-correlation regime under Bayesian learning. Each line reports the portfolio share for households with $q_t \geq 0.8$ minus the corresponding share for households with $q_t \leq 0.2$, in percentage points, equally weighted across households with positive savings.

Figure 6 reports the differences in stock, bond, and cash shares between households with high negative-regime beliefs and those with high positive-regime beliefs. As in the full-information regime-switching model, these differences are relatively small early in the

working life and become substantially larger as households approach retirement. Also consistent with the full-information regime-switching model, households with high negative-regime beliefs hold substantially more long-term government bonds, less stock, and less cash than households with high positive-regime beliefs.²³

Panel B of Table 7 compares these belief-contingent portfolio differences with the corresponding state-contingent differences in the full-information regime-switching benchmark. The two sets of responses are quite similar for households near and during retirement. At age 65, for example, households with high negative-regime beliefs hold approximately 10 percentage points more of their financial portfolio in long-term government bonds than households with high positive-regime beliefs, which is nearly identical to the bond-share gap between the negative- and positive-correlation states under full information. Thus, uncertainty about the current regime only modestly attenuates the portfolio effects generated by persistent variation in stock-bond comovement alone.²⁴ Because observed returns are informative about the latent state, households meaningfully adjust their portfolios in response to changes in the perceived diversification value of long-term government bonds, particularly near and during retirement.

As shown in Figure 6, the magnitude of these belief-contingent portfolio responses varies substantially over the life-cycle. To understand this pattern, Table 8 reports financial wealth as a share of total household resources, defined as beginning-of-period financial wealth plus the expected present value of remaining labor and retirement income, together with the difference in long-term bond shares between high-negative- and high-positive-belief households. Among households with positive savings, financial wealth accounts for only about 3 percent of total resources at age 35 and 8 percent at age 45, but rises to approximately 21 percent at

²³Appendix D shows that these belief-contingent portfolio differences arise primarily through portfolio rebalancing conditional on stock-market participation rather than through changes in the participation margin.

²⁴In Appendix E, we show that these belief-contingent portfolio responses are robust to lowering the per-period fixed cost of holding stocks from 7.3 to 2.0 percent of median household income. At ages 55, 65, and 75, the high-negative-minus-high-positive bond-share gaps are 6.3, 10.3, and 9.6 percentage points under the lower-cost specification, compared with 5.8, 10.0, and 9.4 percentage points in the baseline model.

age 55 and 40 percent by age 65. Over the same ages, the high-negative-minus-high-positive bond-share difference rises from around 1 percentage point to approximately 6 percentage points at age 55 and 10 percentage points by age 65.

Table 8: Life-Cycle Resources and Bond-Share Responses

<i>Panel A: Aggregate life-cycle mechanism</i>						
Age	Financial Wealth Share (%)	Bond-Share Gap (pp)				
35	3.4	1.4				
45	7.8	1.7				
55	20.6	5.8				
65	39.6	10.0				
75	47.8	9.4				
<i>Panel B: Education-specific life-cycle mechanism</i>						
Education	Age 55		Age 65		Age 75	
	Financial Wealth Share (%)	Bond-Share Gap (pp)	Financial Wealth Share (%)	Bond-Share Gap (pp)	Financial Wealth Share (%)	Bond-Share Gap (pp)
High School Dropout	27.2	5.0	45.7	9.1	56.5	7.9
High School Graduate	24.1	6.4	47.8	11.7	55.2	10.2
College Graduate	17.0	5.5	32.1	8.2	39.8	8.5

Notes: Financial wealth is beginning-of-period financial wealth divided by financial wealth plus the expected present value of remaining labor and retirement income. The bond-share gap compares households with $q_t \geq 0.80$ and $q_t \leq 0.20$, among positive savers.

This pattern reflects the changing composition of household resources over the life-cycle. Early in the working life, remaining labor and retirement income represent the majority of household resources, so changes in the covariance structure of financial returns have relatively little effect on the riskiness of the household's overall resource portfolio. As financial wealth accumulates and remaining labor and retirement income decline relative to financial wealth, the composition of the financial portfolio becomes increasingly important for the household's overall risk exposure. As a result, the portfolio effects of changing beliefs about the stock-bond correlation become substantially larger near and during retirement.

Figure 7 shows that the same qualitative pattern holds across education groups. For high school dropouts, high school graduates, and college graduates, high negative-regime beliefs are associated with higher long-term bond shares and lower stock and cash shares, with the differences again becoming most pronounced near and during retirement.




Figure 7: Differences in life-cycle portfolio shares between households with high posterior probability of the negative-correlation regime and households with high posterior probability of the positive-correlation regime, by education. Each panel reports the $q_t \geq 0.8$ minus $q_t \leq 0.2$ differences in stock, long-term bond, and cash shares, in percentage points, equally weighted across households with positive savings.

Taken together, these results show that allowing households to learn about the latent correlation regime preserves the main portfolio implications of the full-information regime-switching benchmark. In particular, households respond to high negative-regime beliefs by shifting financial wealth toward long-term government bonds and away from stocks and cash, with the resulting allocation differences increasing as financial wealth becomes a larger share of total household resources. We next examine which features of the estimated Markov-switching process are most important for determining the magnitude of these belief-contingent portfolio responses.

4.4 The Roles of Persistence and Long-Run Regime Prevalence

To distinguish the roles of regime persistence and long-run regime prevalence in generating the belief-contingent portfolio responses documented in Section 4.3, we parameterize the two-state Markov-switching process in terms of the stationary probability of the negative-correlation regime, $\bar{\pi}_1$, and the second eigenvalue of the transition matrix,

$$\lambda = \pi_{1,1} + \pi_{2,2} - 1.$$

The stationary probability $\bar{\pi}_1$ determines the long-run prevalence of the negative-correlation regime, while λ determines how quickly the influence of current beliefs on the distribution of future regimes decays. This parameterization allows the transition probabilities to be written as

$$\pi_{1,1} = \bar{\pi}_1 + \lambda(1 - \bar{\pi}_1)$$

and

$$\pi_{2,2} = (1 - \bar{\pi}_1) + \lambda\bar{\pi}_1,$$

respectively.

This representation maps directly into the Bayesian learning problem faced by house-

holds. Given posterior beliefs q_t , the probability of the negative-correlation regime h periods ahead satisfies

$$\Pr(z_{t+h} = 1 \mid \mathcal{R}_t) - \bar{\pi}_1 = \lambda^h(q_t - \bar{\pi}_1).$$

The corresponding expected future stock-bond correlation can be written as

$$E_t[\rho_{t+h} \mid \mathcal{R}_t] = \bar{\rho} + \lambda^h(q_t - \bar{\pi}_1)(\rho_1 - \rho_2),$$

where

$$\bar{\rho} = \bar{\pi}_1\rho_1 + (1 - \bar{\pi}_1)\rho_2.$$

When λ is close to one, current beliefs remain informative about expected stock-bond correlations far into the future. When $\lambda = 0$, the rows of the transition matrix are identical, so current beliefs contain no information about the distribution of future regimes beyond the stationary distribution. By contrast, changing $\bar{\pi}_1$ while holding λ fixed changes the long-run prevalence of the two regimes without changing the rate at which the predictive influence of current beliefs decays.

In our baseline model, $\bar{\pi}_1 = 0.166$ and $\lambda = 0.884$. We consider three alternatives, summarized in Table 9. First, we reduce λ to 0.51 while holding $\bar{\pi}_1$ fixed at its baseline value. This substantially reduces regime persistence while preserving the estimated stationary distribution. As shown in Figure 2, the resulting transition matrix, $(\pi_{1,1}, \pi_{2,2}) = (0.591, 0.919)$, lies just inside the low-persistence boundary of the joint 90 percent profile-likelihood region along the constant- $\bar{\pi}_1$ path. We therefore interpret this specification both as a way to assess the role of persistence and as a robustness check that examines how the baseline portfolio results change under a substantially lower but empirically plausible level of persistence. Second, we set $\lambda = 0$ while continuing to hold $\bar{\pi}_1$ fixed at its baseline value. This corresponds to an i.i.d. stationary-mixture economy in which successive correlation regimes are independent and current beliefs contain no predictive information about future diversification opportunities. Finally, we set $\bar{\pi}_1 = 0.5$ while holding λ fixed at its baseline value of 0.884. This last

counterfactual equalizes the long-run prevalence of the two regimes while preserving the rate at which the predictive influence of current beliefs decays.

Table 9: Transition-Process Counterfactuals under Bayesian Learning

Scenario	$\bar{\pi}_1$	λ	$\pi_{1,1}$	$\pi_{2,2}$
Baseline	0.166	0.884	0.903	0.981
Low Persistence	0.166	0.510	0.591	0.919
i.i.d. Regimes	0.166	0.000	0.166	0.834
Equal Long-Run Prevalence	0.500	0.884	0.942	0.942

Notes: $\bar{\pi}_1$ is the stationary probability and long-run prevalence of the negative-correlation regime. $\lambda = \pi_{1,1} + \pi_{2,2} - 1$ is the second eigenvalue and measures persistence: the rate at which the predictive content of current regime beliefs decays. Low Persistence and i.i.d. hold $\bar{\pi}_1$ fixed; Equal Long-Run Prevalence sets $\bar{\pi}_1 = 0.5$ while holding λ fixed.

To isolate the direct effect of the transition process on household portfolio choice, we evaluate each counterfactual using a common distribution of non-belief household states. Specifically, we take the stationary distribution of households across non-belief states from the baseline model and evaluate the corresponding policy rules at posterior beliefs $q_t = 0.8$ and $q_t = 0.2$. We then compute the high-negative-minus-high-positive portfolio-share difference, restricting attention to states in which households choose positive savings under both beliefs. Holding both the non-belief household states and posterior beliefs fixed across counterfactuals allows us to measure how changes in the transition process affect the sensitivity of household portfolio choice rules to a given difference in beliefs.²⁵

Table 10 and Panel (a) of Figure 8 show that reducing regime persistence attenuates the belief-contingent portfolio responses but leaves them economically meaningful under the low-persistence specification. The same basic life-cycle pattern remains. Differences are relatively small early in the working life and become substantially larger near and during retirement. Under the i.i.d. benchmark, by contrast, the common-state response is exactly zero. When

²⁵Because these common-state responses hold other household states fixed and evaluate the policy rules at $q_t = 0.8$ and $q_t = 0.2$, they need not coincide with the simulated conditional averages reported in Section 4.3, which average across the endogenous distribution of household states and posterior beliefs within the $q_t \geq 0.8$ and $q_t \leq 0.2$ groups. Thus, the results in Section 4.3 characterize portfolio differences that arise in the simulated economy, while the common-state exercise we present here isolates how the portfolio choice rules themselves respond to beliefs under alternative transition dynamics.

$\lambda = 0$, the distribution of next period's regime is independent of current beliefs, so the household's portfolio choice rules are invariant to q_t . Changing long-run regime prevalence while holding persistence fixed produces a much smaller change in the belief-contingent portfolio response. Setting $\bar{\pi}_1 = 0.5$ leaves the common-state bond-share response close to its baseline value throughout the life-cycle.




Figure 8: Common-state bond-share responses under alternative transition processes. Each line reports the bond share at $q_t = 0.8$ minus the bond share at $q_t = 0.2$, in percentage points, evaluated using the same distribution of non-belief household states drawn from the baseline model and restricting attention to states that choose positive savings under both beliefs. Panel (a) varies the second eigenvalue λ while holding the stationary probability of the negative-correlation regime fixed at $\bar{\pi}_1 = 0.166$. Panel (b) varies $\bar{\pi}_1$ while holding $\lambda = 0.884$. The low-persistence counterfactual sets $\lambda = 0.51$, just inside the low-persistence boundary of the joint 90 percent profile-likelihood region shown in Figure 2. The i.i.d. case sets $\lambda = 0$.

Across the counterfactuals considered here, belief-contingent portfolio differences are much more sensitive to persistence than to the stationary distribution. This distinction follows directly from the role of λ in the household's forecasting problem. Persistence determines how much information current beliefs contain about future diversification oppor-

tunities, while the stationary distribution primarily determines the long-run frequency with which households encounter the two correlation environments.

Table 10: Transition Dynamics and Common-State Bayesian Portfolio Responses

Scenario	Age 55			Age 65			Age 75		
	Stock	Bond	Cash	Stock	Bond	Cash	Stock	Bond	Cash
Baseline	-2.6	3.4	-0.8	-4.3	6.5	-2.2	-4.2	5.9	-1.7
Low Persistence	-1.3	2.0	-0.8	-2.4	4.2	-1.8	-2.3	3.8	-1.5
i.i.d. Regimes	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Equal Long-Run Prevalence	-2.7	3.3	-0.7	-4.3	6.0	-1.8	-4.1	5.4	-1.3

Notes: Each row evaluates the same distribution of non-belief household states drawn from the baseline Bayesian-learning model at posterior beliefs $q_t = 0.8$ and $q_t = 0.2$. Entries are high-negative minus high-positive portfolio shares among states that choose positive savings under both beliefs. This isolates the direct policy response from cross-sectional composition. At the i.i.d. endpoint ($\lambda = 0$), policy rules are belief invariant and the common-state response is exactly zero.

4.5 Welfare

The preceding results show that persistent variation in the stock-bond correlation can generate meaningful differences in optimal life-cycle portfolio allocations, particularly near and during retirement. We now use the model to quantify the welfare consequences of conditioning household policies on the predictable component of the correlation process.

Specifically, we consider two welfare comparisons. First, we compare the Bayesian-learning policy rules to the policy rules derived under the i.i.d. stationary-mixture economy introduced in Section 4.4 that preserves the stationary mixture of the two correlation regimes and therefore the unconditional return environment, while eliminating the predictive content of current beliefs for future diversification opportunities. Both sets of policy rules are evaluated along the same realized household shocks and return histories generated by the estimated Markov-switching process. The resulting consumption-equivalent gain therefore measures the welfare cost of using policy rules that ignore the predictable component of the correlation process. Next, we compare the full-information policy rules with the Bayesian-learning policy rules along the same realized household shocks and return histories. This second comparison measures the additional value of observing the current latent correlation

regime perfectly rather than inferring it from observed returns.

We report welfare in consumption-equivalent units using the complete household objective, including utility from consumption and bequests. Because the bequest term is not homogeneous in consumption, we compute the consumption-equivalent compensation numerically as follows. Let $\bar{U}_C^B$ and $\bar{U}_B^B$ denote the average lifetime utility contributions from consumption and bequests under the benchmark policy rules and let $\bar{U}^A$ denote average lifetime utility under the alternative policy rules. The resulting consumption-equivalent variation (CEV), denoted by κ , satisfies

$$(1 + \kappa)^{1-\gamma} \bar{U}_C^B + \bar{U}_B^B = \bar{U}^A \quad (19)$$

and measures the proportional increase in consumption under the benchmark policy rules required to make households indifferent between the two alternatives, while holding fixed the bequest utility generated by the benchmark policy.

Table 11: Welfare Effects under Bayesian Learning

Group	Bayesian vs. i.i.d. (%)	Full Information vs. Bayesian (%)
Overall	0.0044	0.0042
High School Dropout	0.0046	0.0052
High School Graduate	0.0046	0.0039
College Graduate	0.0016	0.0013

Notes: Consumption-equivalent gains use the complete household objective, including warm-glow bequests. Bayesian vs. i.i.d. measures the value of exploiting persistent covariance predictability with posterior-belief-contingent policies. Full Information vs. Bayesian measures the additional value of observing the current latent correlation regime perfectly.

Table 11 reports our results. The aggregate welfare gain from replacing policy rules that ignore covariance predictability with the Bayesian-learning model's belief-contingent rules is positive but quantitatively modest, at just 0.0044 percent, or approximately 0.44 basis points of lifetime consumption. The gain is largest for high school dropouts and high school graduates and substantially smaller for college graduates. This pattern is consistent with Table 8, where financial wealth represents a smaller share of total resources for college

graduates at the ages when belief-contingent portfolio responses are largest.

Perfect observation of the current correlation regime provides an additional welfare gain of 0.0042 percent, or approximately 0.42 basis points of lifetime consumption, relative to the Bayesian-learning model. The gains are again largest for high school dropouts and high school graduates and smaller for college graduates. Thus, uncertainty about the latent stock-bond correlation regime imposes a positive but modest welfare cost, although the value of eliminating that uncertainty remains modest in consumption-equivalent terms.²⁶

The small aggregate welfare gains are notable given the sizable portfolio differences documented above. The portfolio-share results characterize how the composition of optimal portfolios varies with beliefs about the correlation regime, whereas the consumption-equivalent variation measures the utility cost of following an alternative policy rule that remains reasonably well matched to the return environment over much of the relevant state space.

Table 12: Distribution of Realized High-Negative-Belief Exposure

Exposure Group	Share of Households (%)	Mean Years	Mean Share of Lifetime (%)
0 years	45.6	0.0	0.0
1–5 years	22.5	2.8	4.8
6–10 years	14.2	7.8	13.1
11–20 years	13.0	14.6	24.3
21+ years	4.6	26.7	43.6

Notes: Exposure is the realized number of years in which the posterior probability of the negative-correlation regime satisfies $q_t \geq 0.80$. Share of Households is the percentage of the 10 million simulated household lifetimes belonging to each exposure group. Mean Share of Lifetime is the within-group average fraction of each household's realized lifetime spent with $q_t \geq 0.80$.

To better understand why economically meaningful portfolio reallocations coexist with small average welfare differences, we next examine how welfare gains vary with realized exposure to high negative-regime beliefs. Table 12 groups households according to the realized number of years in which the posterior probability of the negative-correlation regime satis-

²⁶The two consumption-equivalent gains are computed separately and are not mechanically additive because the bequest component is not homogeneous in consumption.

fies $q_t \geq 0.8$ and Figure 9 displays the corresponding consumption-equivalent welfare gains.²⁷

The welfare gain from the Bayesian-learning model relative to the i.i.d. stationary-mixture model rises sharply with high-negative-belief exposure. Among households that never experience a high-negative belief, the gain is only 0.0006 percent, or approximately 0.06 basis points of lifetime consumption. It rises to 0.0062 percent for households with six to ten years of exposure, 0.013 percent for those with eleven to twenty years, and 0.029 percent, or approximately 2.9 basis points, for households with more than twenty years of exposure.




Figure 9: The solid black line reports consumption-equivalent gains from Bayesian learning relative to the i.i.d. stationary-mixture benchmark, and the dashed red line reports the additional gains from full information relative to Bayesian learning. Exposure groups are defined by the realized number of lifetime years in which the posterior probability of the negative-correlation regime satisfies $q_t \geq 0.80$.

The additional value of full information also increases with high-negative-belief exposure, as shown in Figure 9. Households that spend little or no time with a high posterior

²⁷Because these groups are defined using realized future belief histories, the exercise is a mechanism decomposition rather than an alternative ex ante welfare criterion.

probability of the negative-correlation regime therefore have relatively little to gain from exploiting persistence in the correlation process. By contrast, households that repeatedly encounter periods in which bonds are perceived to provide strong diversification benefits face larger cumulative welfare costs from using policy rules that ignore covariance predictability. Thus, average welfare gains are small because most households spend relatively little of their lives with a high posterior probability of the negative-correlation regime, while the gains rise substantially for households that experience such periods repeatedly.

5 Conclusion

We find that persistent variation in the stock-bond correlation can materially affect optimal life-cycle portfolio choice even when expected returns and marginal volatilities are held fixed across regimes. Households with a high posterior probability of the negative-correlation regime shift financial wealth toward long-term government bonds and away from stocks and cash, with these responses becoming substantially larger near and during retirement. The magnitude of the response depends primarily on regime persistence rather than long-run regime prevalence. Despite these sizable allocation effects, the average welfare gain from exploiting covariance predictability is modest, at approximately 0.4 basis points of lifetime consumption, although the gains are considerably larger for households with sustained high-negative-belief exposure.

More broadly, our results suggest that persistent covariance variation can be an economically relevant state variable in quantitative models of portfolio choice over the life-cycle even when expected returns and marginal volatilities are constant. For the portfolio responses we study, the persistence of the covariance process is substantially more important than the long-run prevalence of the underlying regimes because persistence determines how informative current beliefs are about future diversification opportunities. This channel becomes more important as financial wealth rises relative to remaining labor and retirement income,

helping explain why the portfolio effects are concentrated near and during retirement. At the same time, the modest average welfare gains show that sizable state-contingent portfolio reallocations need not translate into large ex ante welfare effects.

Several limitations of our study suggest potentially fruitful directions for future work. First, our two-state Markov specification, motivated by the limited length of the annual returns sample, is a simplification relative to richer regime-switching models of the joint stock-bond distribution, such as the four-regime framework of [Guidolin and Timmermann \(2007\)](#). A longer sample or higher-frequency data would allow for sharper identification of a richer state space while mitigating the identification concerns that arise in our relatively short annual sample. Extending the analysis to international data could provide an additional source of variation for identifying the correlation process.

Second, households in our model take the parameters governing the return process as given and learn only about the latent correlation regime. Allowing households to learn jointly about the latent state and the parameters governing the return process would introduce an additional source of uncertainty that could affect both the speed of belief updating and the resulting portfolio effects. This would substantially increase the dimensionality of the household problem, but it represents a natural extension of the Bayesian-learning framework considered here.

Our model is also a partial-equilibrium model that treats asset returns as exogenous. This omission could matter if belief-contingent portfolio demand were large enough to affect equilibrium asset prices and, hence, the underlying stock-bond correlation process itself.

Finally, our asset menu is restricted to cash, stocks, and long-term government bonds. Incorporating housing, annuities, tax-advantaged retirement accounts, and institutional frictions such as target-date fund constraints or limited rebalancing opportunities would be natural but potentially computationally demanding extensions.

Declaration of Generative AI Use

The authors used generative AI tools, including ChatGPT, Claude and Refine.ink, during the preparation of this manuscript for proofreading, language editing, and formatting assistance, as well as for brainstorming potential referee critiques to help identify areas of the manuscript that could benefit from clarification. All modeling decisions, calibration choices, computational implementation, and substantive analysis were conducted, reviewed, and verified by the authors. The authors take full responsibility for the content of this paper, including any errors or omissions.

6 References

- Ang, A. and Bekaert, G. (2002). International asset allocation with regime shifts. *The Review of Financial Studies*, 15(4):1137–1187.
- Arias, E. and Xu, J. (2022). United States Life Tables, 2019. *National Vital Statistics Reports*, 71(1).
- Ashman, H. and Neumuller, S. (2020). Can income differences explain the racial wealth gap? A quantitative analysis. *Review of Economic Dynamics*, 35:220–239.
- Baele, L., Bekaert, G., and Inghelbrecht, K. (2010). The determinants of stock and bond return comovements. *The Review of Financial Studies*, 23(6):2374–2428.
- Benzoni, L., Collin-Dufresne, P., and Goldstein, R. S. (2007). Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *The Journal of Finance*, 62(5):2123–2167.
- Campbell, J. Y., Pflueger, C., and Viceira, L. M. (2020). Macroeconomic drivers of bond and equity risks. *Journal of Political Economy*, 128(8):3148–3185.

- Campbell, J. Y. and Viceira, L. M. (1999). Consumption and portfolio decisions when expected returns are time varying. *The Quarterly Journal of Economics*, 114(2):433–495.
- Campbell, J. Y. and Viceira, L. M. (2002). *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economics.
- Cocco, J. F., Gomes, F. J., and Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *The Review of Financial Studies*, 18(2):491–533.
- Connolly, R., Stivers, C., and Sun, L. (2005). Stock market uncertainty and the stock-bond return relation. *Journal of Financial and Quantitative Analysis*, 40(1):161–194.
- Damodaran, A. (2026). Historical Returns on Stocks, Bonds and Bills: 1928-2025. http://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/histretSP.html. Accessed: 2026-02-22.
- De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies*, 71(3):743–768.
- De Nardi, M. and Yang, F. (2014). Bequests and heterogeneity in retirement wealth. *European Economic Review*, 72:182–196.
- Fagereng, A., Gottlieb, C., and Guiso, L. (2017). Asset market participation and portfolio choice over the life-cycle. *The Journal of Finance*, 72(2):705–750.
- Gomes, F. and Michaelides, A. (2005). Optimal life-cycle asset allocation: Understanding the empirical evidence. *The Journal of Finance*, 60(2):869–904.
- Guidolin, M. and Timmermann, A. (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11):3503–3544.
- Ilmanen, A. (2003). Stock-bond correlations. *The Journal of Fixed Income*, 13(2):55–66.

- Kandel, S. and Stambaugh, R. F. (1996). On the predictability of stock returns: an asset-allocation perspective. *The Journal of Finance*, 51(2):385–424.
- Kotlikoff, L. J. and Summers, L. H. (1981). The role of intergenerational transfers in aggregate capital accumulation. *Journal of Political Economy*, 89(4):706–732.
- Liu, H. (2011). Dynamic portfolio choice under ambiguity and regime switching mean returns. *Journal of Economic Dynamics and Control*, 35(4):623–640.
- Lockwood, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, 108(9):2513–2550.
- Molenaar, R., Sénéchal, E., Swinkels, L., and Wang, Z. (2024). Empirical evidence on the stock–bond correlation. *Financial Analysts Journal*, 80(3):17–36.
- Semega, J. L., Fontenot, K. R., and Kollar, M. A. (2017). Income and poverty in the United States: 2016. Current Population Reports P60-259, U.S. Census Bureau.
- Tischbirek, A. (2019). Long-term government debt and household portfolio composition. *Quantitative Economics*, 10(3):1109–1151.
- 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.

Appendix

A Computational Implementation

This appendix describes the numerical implementation of the model. We first summarize the state space and numerical representation of the household problem, then present the backward-induction algorithm used to solve the Bayesian-learning model. We conclude by describing the simulation and welfare procedures used to construct the quantitative results.

A.1 State Space and Numerical Representation

State space. At age t , a working household's decision problem is indexed by financial wealth a_t , the persistent and transitory components of income (ψ_t, e_t) , education j , discount-factor type β_i , and the posterior probability q_t of the negative-correlation regime. During retirement, the transitory income state drops out and the persistent component is fixed at its retirement value ψ_{65} .

Discretization. Financial wealth is represented on a nonuniform grid of 400 points on $[-8, 57]$, expressed relative to median household income, with additional resolution around zero. For each education group, the persistent income process is approximated with five states and the transitory income shock with two equally likely realizations. Stock and bond portfolio shares are discretized in increments of 0.02. Together with the no-short-selling and no-leverage restrictions in equations (9)–(12), this yields 1,326 feasible stock-bond-cash portfolios. The posterior belief q_t is represented on 11 equally spaced points on $[0, 1]$.

Return integration and belief updating. For each correlation regime, we approximate the regime-conditional distribution of stock and bond returns using 100 equally weighted return realizations. The selected nodes reproduce the target means and standard deviations to within 0.001 and the target stock-bond correlation to within 0.01. The same return support is used in the Bayesian-learning model, the matched full-information benchmark, and the

counterfactual exercises.

For each current belief q_t , we first compute the predictive probability $\tilde{q}_{t+1}$ using equation (7). For each regime-specific return realization, we then evaluate the continuous Gaussian likelihoods and update the posterior using equation (8). The resulting posterior is mapped to the two adjacent points on the belief grid using linear interpolation. Because a realized return pair changes both next-period wealth and the posterior belief, the Bayesian solver combines these mappings in sparse joint wealth-belief transition operators.

A.2 Solution Algorithm

We solve the finite-horizon household problem by backward induction from age 100 to age 25. The numerical procedure can be summarized as follows.

Algorithm 1: Bayesian-Learning Solution.

1. **Construct transition objects.** For each regime, savings choice, feasible portfolio, and return node, compute next-period wealth and map it onto the financial-wealth grid. Combine these wealth transitions with the posterior-belief interpolation weights described above to construct the sparse joint wealth-belief transition operators.
2. **Initialize the terminal problem.** At age 100, initialize continuation values using the warm-glow bequest function.
3. **Solve the retirement problem.** For $t = 100, \dots, 65$, iterate over education, persistent-income, discount-factor, wealth, and belief states. For each savings choice, apply the joint wealth-belief transition operator, combine survival and bequest utility with the continuation value, and maximize over feasible portfolios. Add current-period utility and maximize over savings to obtain the age- t value, savings, and portfolio policy functions.
4. **Solve the working-life problem.** For $t = 64, \dots, 25$, first integrate next-period values over the persistent- and transitory-income transitions. For each household state and sav-

ings choice, apply the joint wealth-belief transition operator and maximize over feasible portfolios. Add current-period utility and maximize over savings to obtain the corresponding policy functions.

5. **Solve the matched full-information benchmark.** Repeat the same backward-induction procedure using the same household grids and return integration nodes, but index continuation values and policy rules by the observed latent correlation regime z_t rather than by posterior beliefs.

The Bayesian and full-information solutions therefore use the same return process, household-state discretization, portfolio grid, and numerical integration. The difference between the two specifications is the information available to households.

A.3 Simulation and Welfare

Simulation. We construct the stationary distribution by simulating 10 million household lineages for 2,000 annual periods. A lineage continues when one household is replaced by its age-25 descendant. We allow the latent return process and the posterior belief implied by the stock and bond return history to continue across generations. For each lineage and simulation period, we perform the following steps:

1. Draw the innovations governing the regime transition, return realization, survival, education transmission, delayed inheritance, and income shocks.
2. Transition the latent correlation regime and draw stock and bond returns conditional on that regime.
3. Apply the realized return to the preceding savings and portfolio choices, after the first simulation period, to obtain beginning-of-period financial wealth, and update the posterior belief.

4. Apply survival and aging. If the current household dies, transfer the estate according to the intergenerational rules, initialize the age-25 descendant, and carry forward the latent regime and posterior belief.
5. Realize any delayed inheritance and update the persistent and transitory income states.
6. Evaluate the household's saving policy and, conditional on positive savings, its portfolio policy.

After 2,000 periods, the cross-sectional distribution is used to construct the wealth, participation, portfolio, and posterior-belief statistics reported in the paper. Household-specific random-number seeds are held fixed across information structures so that comparisons are evaluated using the same exogenous shock histories.

Welfare. The welfare exercises draw 10 million age-25 entrants, with replacement, from the stationary Bayesian cross-sectional distribution. For each entrant, the Bayesian, i.i.d. stationary-mixture, and full-information policy rules are evaluated along identical realizations of correlation states, asset returns, income shocks, mortality, and delayed inheritances. Under each policy specification, the household is followed until death or age 100, discounted consumption utility and warm-glow bequest utility are accumulated, and average lifetime utility is computed across entrants. The Bayesian-versus-i.i.d. and full-information-versus-Bayesian consumption-equivalent gains are then obtained by solving equation (19) numerically. All welfare policy rules are evaluated under the estimated persistent Markov-switching return process. The policy rule being applied differs across comparisons, while the full-information comparison additionally changes the information available to households.

B Bayesian Numerical Accuracy

The Bayesian-learning model includes the posterior probability of the negative-correlation regime, q_t , as a state variable in the household problem. We approximate this state using a finite belief grid. The baseline results in the main text use 11 belief points together with the same 100 regime-specific return nodes used in the full-information benchmark. This appendix examines the sensitivity of the belief-contingent portfolio responses to refinement of the belief grid.

Table 13 reports the high-negative-minus-high-positive bond-share differences at ages 45, 55, 65, and 75 using 7, 9, and 11 belief points. In each case, the return discretization is held fixed at 100 nodes per regime. The resulting portfolio responses are stable across the alternative grids, particularly at the ages when the belief-contingent portfolio differences are largest.

Table 13: Bayesian Belief-Grid Convergence

Belief Points	Return Nodes	Age 45	Age 55	Age 65	Age 75
7	100	1.75	5.69	10.05	9.35
9	100	1.69	5.78	9.97	9.16
11	100	1.72	5.80	9.96	9.37

Notes: Entries are simulated high-negative minus high-positive belief bond-share gaps in percentage points. All specifications use the original 100 return nodes per regime; only the posterior-belief grid is refined.

The return nodes are used only to evaluate expectations in the household problem. Bayesian updating is based on the continuous Gaussian regime-specific densities implied by the estimated Markov-switching model. Thus, observing a particular numerical return node does not mechanically reveal the underlying latent state. This allows the Bayesian and full-information models to use the same numerical return integration while differing only in the information available to households.

C Sensitivity to Regime-Dependent Expected Returns

The baseline specification in the main text restricts stock and bond means and marginal volatilities to be constant across regimes, allowing only the stock-bond correlation and transition probabilities to vary by state. In this appendix, we examine the sensitivity of both the estimated regime process and the resulting portfolio allocations to a relaxation of the fixed-mean restriction. Specifically, we estimate a partially unrestricted Markov-switching specification in which stock and bond means, the stock-bond correlation, and transition probabilities are allowed to vary across regimes, while stock and bond marginal volatilities remain common across states. We then re-solve and re-simulate the Bayesian-learning model under the partially unrestricted estimates, holding all remaining calibrated parameters fixed at their baseline values.

Table 14 compares the resulting estimates to those from the baseline restricted MLE. The estimated regime-dependent correlations and transition probabilities are very similar across the two specifications. Under the partially unrestricted specification, the negative-correlation regime has $\rho_1 = -0.82$ and $\pi_{1,1} = 0.905$, compared with $\rho_1 = -0.81$ and $\pi_{1,1} = 0.903$ in the baseline restricted MLE. Similarly, the positive-correlation regime has $\rho_2 = 0.25$ and $\pi_{2,2} = 0.984$, compared with $\rho_2 = 0.25$ and $\pi_{2,2} = 0.981$ in the restricted specification.

The primary difference between the two specifications lies in the estimated state-dependent means. Under the partially unrestricted MLE, the negative-correlation regime has a lower estimated stock mean and a higher estimated bond mean, $(\mu_{s,1}, \mu_{b,1}) = (0.040, 0.031)$, while the positive-correlation regime has a higher estimated stock mean and lower estimated bond mean, $(\mu_{s,2}, \mu_{b,2}) = (0.095, 0.016)$. These estimates should be interpreted with caution, however, given the relatively short annual sample and the difficulty of estimating conditional expected returns. Importantly, allowing the means to vary leaves the estimated regime-dependent correlations, transition probabilities, and implied regime timing nearly unchanged. This suggests that the baseline regime classification is not driven mechanically by the fixed-mean restriction.

Table 14: Sensitivity to Regime-Dependent Expected Returns

Panel A: Return-process estimates

Parameter	Restricted MLE	Partially Unrestricted MLE
State 1 stock mean	0.0861	0.0402
State 1 bond mean	0.0183	0.0312
Stock volatility	0.1930	0.1909
Bond volatility	0.0886	0.0886
State 1 correlation	-0.8136	-0.8215
State 1 persistence	0.9030	0.9053
State 2 stock mean	0.0861	0.0953
State 2 bond mean	0.0183	0.0158
State 2 correlation	0.2518	0.2506
State 2 persistence	0.9807	0.9836
Log likelihood	127.1984	127.7491

Panel B: Bayesian portfolio-share responses

	Restricted Baseline				Partially Unrestricted		
Age	Stock	Bond	Cash		Stock	Bond	Cash
55	-4.37	5.80	-1.43		-68.93	70.96	-2.03
65	-6.46	9.96	-3.50		-62.43	66.73	-4.30
75	-6.34	9.37	-3.03		-56.61	60.95	-4.34

Panel C: Posterior diagnostics

Diagnostic	Restricted Baseline	Partially Unrestricted
Posterior confidence ≥ 0.80 (%)	85.7	88.3
Latent-state classification accuracy (%)	91.8	93.1
Brier score	0.063	0.054

Notes: Panel A compares the restricted baseline MLE with the partially unrestricted MLE, which allows stock and bond means to vary across regimes and estimates common marginal volatilities jointly with the regime correlations and transition probabilities. Panel B re-solves and re-simulates the Bayesian-learning model under each estimated return process while holding all household-side calibrated parameters fixed. Entries are high-negative minus high-positive posterior-belief portfolio-share differences, in percentage points, among households with positive savings; high-negative and high-positive beliefs satisfy $q_t \geq 0.80$ and $q_t \leq 0.20$. Panel C reports selected posterior-informativeness diagnostics from the corresponding simulations.

To examine the quantitative implications of these differences, we re-solve and re-simulate the Bayesian-learning model using the partially unrestricted return-process estimates. Panel B of Table 14 compares the resulting belief-contingent portfolio responses with those from the baseline model. Allowing expected returns to vary across regimes substantially increases the magnitude of the reallocation between stocks and long-term government bonds. At age 65, the high-negative-minus-high-positive posterior-belief bond-share difference increases from 10.0 percentage points in the baseline model to 66.7 percentage points under the partially unrestricted specification, while the corresponding stock-share difference increases in magnitude from -6.5 to -62.4 percentage points. The same pattern is evident at ages 55 and

75, while the corresponding differences in cash holdings remain comparatively small.

These larger portfolio responses reflect the additional information about expected returns contained in beliefs about the latent regime. In the negative-correlation regime, the lower expected stock return and higher expected bond return reinforce the greater diversification value of long-term government bonds. In the positive-correlation regime, the higher expected stock return and lower expected bond return strengthen the reallocation toward stocks. As a result, the partially unrestricted specification generates substantially larger belief-contingent portfolio differences than the baseline model, in which beliefs affect portfolio choice only through variation in stock-bond covariance.

Panel C shows that the additional state dependence in expected returns also makes the latent regime somewhat easier to infer from realized returns. The fraction of simulated observations with posterior confidence of at least 0.80 increases from 85.7 to 88.3 percent, while latent-state classification accuracy rises from 91.8 to 93.1 percent and the Brier score falls from 0.063 to 0.054. These relatively modest changes in posterior informativeness suggest that the much larger portfolio responses primarily reflect the additional expected-return channel rather than a large change in the precision with which households are able to infer the latent state.

Figure 10 shows that these differences persist across the life cycle. The partially unrestricted specification produces substantially larger shifts from stocks toward long-term government bonds as posterior beliefs move toward the negative-correlation regime, while the cash response remains comparatively small.

These results provide an additional rationale for retaining the restricted specification as the baseline quantitative experiment. The estimated correlations and transition dynamics remain nearly unchanged when expected returns are allowed to vary across regimes, and the direction of the belief-contingent portfolio response is preserved. The substantially larger portfolio differences under the partially unrestricted specification arise because beliefs in this case affect both expected returns and diversification opportunities.




Figure 10: Belief-contingent portfolio-share responses under the restricted baseline and partially unrestricted return processes. Each panel reports the high-negative-minus-high-positive posterior-belief portfolio-share difference among households with positive savings. High-negative and high-positive beliefs satisfy $q_t \geq 0.8$ and $q_t \leq 0.2$, respectively. The partially unrestricted specification allows stock and bond expected returns, the stock-bond correlation, and transition probabilities to vary across regimes while maintaining common marginal volatilities. All household-side calibrated parameters are held fixed at their baseline values. Portfolio-share differences are reported in percentage points.

D Stock-Market Participation and Rebalancing

The portfolio-share results in the main text are computed among all households with positive savings and therefore combine differences in stock-market participation with differences in portfolio composition conditional on participation.

To determine which margin drives the belief-contingent portfolio responses reported in Section 4.3, Figure 11 examines these margins separately. The participation profiles in Panel (a) of Figure 11 are relatively similar across the two belief groups, indicating that changes in beliefs about the correlation regime have little effect on stock-market participation rates. By contrast, Panel (b) shows economically meaningful portfolio rebalancing among stock-market participants as households with high negative-regime beliefs shift financial wealth toward long-term bonds and away from stocks and cash. Among stock-market nonparticipants, the corresponding reallocation between bonds and cash in Panel (c) is comparatively small. Thus, the belief-contingent portfolio differences reported in the main text primarily reflect changes in portfolio composition among households that participate in the stock market rather than changes in the extensive participation margin or a reallocation between bonds and cash for households that do not hold stock.




Figure 11: Stock-market participation and portfolio rebalancing under Bayesian learning. Panel (a) compares stock-market participation among households with $q_t \geq 0.8$ and households with $q_t \leq 0.2$, conditional on positive savings. Panel (b) reports high-negative-minus-high-positive posterior-belief differences in stock, long-term bond, and cash shares among stock-market participants. Panel (c) reports the corresponding differences in long-term bond and cash shares among stock-market nonparticipants. Portfolio-share differences in Panels (b) and (c) are reported in percentage points.

E Robustness to the Fixed Cost of Holding Stocks

The baseline calibration sets the per-period fixed cost of holding stocks at $\zeta = 0.073$, or 7.3 percent of median household income. This appendix examines whether the main belief-contingent portfolio results are sensitive to this choice by re-solving and re-simulating the Bayesian-learning model with $\zeta = 0.020$, holding all other calibrated parameters and the estimated return process fixed at their baseline values.

Table 15 compares the high-negative-minus-high-positive posterior-belief portfolio-share differences in the baseline model and under the lower fixed cost specification. The belief-contingent portfolio responses are very similar across the two calibrations. At ages 55, 65, and 75, the long-term bond-share differences are 6.3, 10.3, and 9.6 percentage points under the lower fixed cost, compared with 5.8, 10.0, and 9.4 percentage points in the baseline model. The corresponding stock and cash responses also remain similar. Thus, lowering the fixed cost of holding stocks substantially below its baseline calibrated value does not materially alter the portfolio response to beliefs about the stock-bond correlation regime.

Table 15: Lower Stock Fixed-Cost Sensitivity under Bayesian Learning

Specification	Stock-Share Gap			Bond-Share Gap			Cash-Share Gap		
	55	65	75	55	65	75	55	65	75
Baseline ($\zeta = 0.073$)	-4.37	-6.46	-6.34	5.80	9.96	9.37	-1.43	-3.50	-3.03
Lower Cost ($\zeta = 0.020$)	-4.82	-6.73	-6.64	6.29	10.32	9.62	-1.47	-3.59	-2.99

Notes: Entries are high-negative minus high-positive posterior-belief portfolio-share differences, in percentage points, among households with positive savings. High-negative and high-positive beliefs satisfy $q_t \geq 0.80$ and $q_t \leq 0.20$, respectively. The lower-cost specification sets the per-period fixed cost of holding stocks to $\zeta = 0.020$ while holding all other calibrated parameters, the estimated return process, and the Bayesian learning specification fixed at their baseline values.

Stock-Bond Correlation Regimes and Portfolio Choice over the Life-Cycle*

Seth Neumuller[†]

Jess Yao[‡]

Wellesley College

Wellesley College

September 2026

Abstract

The stock-bond correlation varies substantially over time, altering the diversification benefits of long-term government bonds. We embed a two-state Markov-switching process for the stock-bond correlation into an incomplete-markets life-cycle model in which households invest in cash, stocks, and long-term government bonds. Households infer the latent correlation regime from observed returns and condition their portfolios on posterior beliefs. The model generates sizable differences in optimal portfolios across beliefs, with persistence in the estimated correlation process playing a central role in their magnitude. These effects are concentrated near and during retirement, as financial wealth becomes a larger share of household resources. Despite these sizable allocation effects, conditioning portfolios on covariance predictability raises ex ante welfare by only 0.4 basis points of lifetime consumption on average. However, welfare gains are substantially larger for households that experience sustained periods of high posterior probability of the negative-correlation regime, reaching approximately 2.9 basis points for households with more than twenty years of such exposure.

Keywords: Life-Cycle Portfolio Choice; Stock-Bond Correlation; Bayesian Learning; Incomplete Markets

JEL Codes: G11, D91, G12, E21

*We thank seminar participants at Wellesley College, Middlebury College, Amherst College, and UMass Lowell for helpful comments and suggestions. All remaining errors and omissions are our own.

[†]Corresponding Author; seth.neumuller@wellesley.edu; PNE 423, Department of Economics, Wellesley College, 106 Central St., Wellesley, MA 02481

[‡]zy102@wellesley.edu