FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle? Redux

Gleb Kozliakov
University of California, Davis

Emile A. Marin
University of California, Davis

Sanjay R. Singh
Federal Reserve Bank of San Francisco
University of California, Davis

March 2026

Working Paper 2026-06

https://doi.org/10.24148/wp2026-06

Suggested citation:

Kozliakov, Gleb, Emile A. Marin, and Sanjay R. Singh. 2026. "Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle? Redux." Federal Reserve Bank of San Francisco Working Paper 2026-06. https://doi.org/10.24148/wp2026-06

Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle? Redux*

Gleb Kozliakov† UC Davis

Emile A. Marin‡ UC Davis

Sanjay R. Singh§ FRBSF and UC Davis

March 2026

Abstract

Can idiosyncratic risk explain the equity premium? We revisit this question using a novel measure of imperfect risk sharing, implied by a large class of heterogeneous-agent models, constructed using household-level panel data. We identify a group of households – with relatively high income but low net-worth – whose consumption is sufficiently volatile and risky to explain 94% of the observed U.S. Sharpe ratio. In contrast, the consumption dynamics of high net-worth individuals predict a negative Sharpe ratio and so do not constitute the relevant pricing factor, consistent with models featuring wealth motives.

Keywords: uninsurable idiosyncratic risk, heterogeneous agents, wealth dynamics, equity premium

JEL codes: G12, B52, E21

*We thank Luigi Bocola, Nicolas Caramp, Keshav Dogra, Tarek Hassan, Ben Hébert, Zhengyang Jiang, Arvind Krishnamurthy, Thomas Mertens, Lawrence Schmidt, and participants at various conferences and seminars for useful feedback. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. Remaining errors are our own.

†Email: gkozliakov@ucdavis.edu. Address: One Shields Avenue, Davis, CA, 95616, USA.

‡Email: emarin@ucdavis.edu. Address: One Shields Avenue, Davis, CA, 95616, USA.

§Email: Sanjay.Singh@sf.frb.org. Address: 101 Market Street, San Francisco, CA 94105, USA.

1 Introduction

A classical literature investigates whether idiosyncratic risk and imperfect risk sharing across households (Constantinides and Duffie, 1996) can resolve the equity premium puzzle (Mehra and Prescott, 1985). Lettau (2002) extends the Hansen and Jagannathan (1991) volatility bounds and, using estimates of idiosyncratic income risk within a no-trade asset pricing framework, concludes that existing risk measures exhibited weak cyclical properties and could explain at best about 20% of the observed U.S. equity market Sharpe ratio, for an elasticity of intertemporal substitution (EIS) of 0.2. In this paper, we revisit these calculations using a novel measure of imperfect risk sharing based on an "as-if" representative agent Euler equation. For the same EIS of 0.2, we find that this measure can explain approximately 94% of the historical U.S. Sharpe ratio.

In a large class of incomplete-markets models, the aggregate implications of imperfect risk sharing can be summarized by a discount-factor ($\beta-$wedge) in the aggregate Euler equation (Nakajima, 2005). Building on Berger, Bocola, and Dovis (2023), we measure this wedge directly from household-level panel data on consumption, income, and assets available in the Consumption Expenditure Survey (CEX), assuming power utility over consumption for a range of EIS values. Consistent with theory, we identify the marginal "pricer" of risk-free bonds as households with above-median income but below-median wealth within their income category (henceforth *HYLW* households). These agents are currently unconstrained yet highly exposed to the risk of hitting their borrowing limits, and so are willing to pay the highest price for risk-free bonds. The $\beta-$wedge constructed from this group of households is both volatile and countercyclical, giving it the potential to explain asset pricing puzzles.

With this new measure of imperfect risk-sharing, we revisit the empirical exercise in Lettau (2002) and derive the model-implied Sharpe ratio under the assumption that the pricers of risk-free bonds also price equity. The Sharpe ratio admits a simple representation:$(\gamma - \phi)\sigma_\varepsilon$, where ϕ is the cyclicality of the $\beta-$wedge with respect to aggregate consumption growth and σ_ε is the volatility of aggregate consumption shock. A negative ϕ implies the marginal

Figure 1. Baseline Results

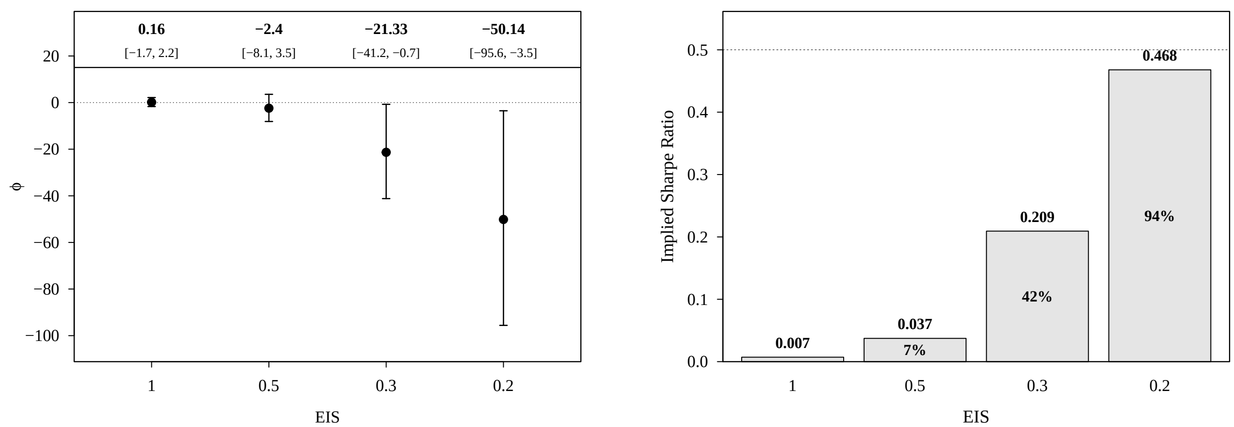

(a) Countercyclicality Parameter ϕ **(b) Model-implied Sharpe ratio**

Notes. Figure 1a plots the estimated countercyclicality parameter ϕ across values of the elasticity of intertemporal substitution (EIS). The 95% confidence intervals are constructed using a parametric AR(1) bootstrap procedure, see Appendix A.4 for details. Figure 1b reports the model-implied Sharpe ratio for each EIS value. The horizontal dashed line denotes the benchmark historical Sharpe ratio of 0.5. Black labels inside the bars indicate the percentage of the benchmark Sharpe ratio explained by the model. In all estimations, the aggregate shock volatility is calibrated to $\sigma_\varepsilon = 0.8\%$ such that volatility of aggregate consumption growth is 1.1% .

investor faces greater uninsurable consumption risk during recessions. The role of the EIS (γ^{-1}) is economically meaningful and more than a simple scale factor: as utility becomes more concave, the $\beta-$wedge places greater weight on households within the HYLW group experiencing low consumption growth. Figure 1 illustrates the interaction between the cyclicality of the $\beta-$wedge and the EIS for explaining the Sharpe ratio on equity. For an EIS of 0.2, the $\beta-$wedge is strongly countercyclical and the model implies a Sharpe ratio of 0.47, or 94% of the value observed in the data.[1]

We investigate the mechanisms driving these large implied Sharpe ratios. The β-wedge is constructed as the cross-sectional average of inverse consumption-share changes within a specified group of households. As such, this encodes both a group mean and a within-group dispersion term. The group mean component captures movements in the group's average

[1]Further decreasing the EIS increases the model-implied estimates. Explaining more than 100% of the observed Sharpe ratio is desirable because, as Bansal and Lehmann (1997) and others note, estimates based on the S&P 500 likely underestimate the true U.S. market Sharpe ratio.

consumption relative to aggregate consumption, whereas the dispersion term (a Jensen's term) isolates idiosyncratic risk. Indeed, we find that within-group dispersion drives the countercyclicality of the measured wedge and hence premia, rather than variation in the average group term. This supports the interpretation that idiosyncratic risk within the HYLW group is behind the large equity premia predicted by the model. Relatedly, we split the sample into periods corresponding to good and bad aggregate states and show that, while the $\beta-$wedge is countercyclical in both subsamples, substantially higher average Sharpe ratios arise during downturns. Through the lens of the model, HYLW households face greater consumption risk in bad times.

In contrast, wedges constructed from data on high-income, high-wealth households (henceforth HYHW) are procyclical and generate negative Sharpe ratios, implying that their consumption dynamics cannot account for positive equity premia. These households constitute over 40% of all assetholders in the CEX vis-a-vis 20% which are HYLW.[2] How, then, do we reconcile that only the consumption dynamics of the latter group appear able to price risky assets? Using a stylized model with wealth in the utility function, we show that both HYLW and HYHW households agree on asset prices (bonds and equity) and yet consumption growth is a poor proxy for the latter group's valuations.

Furthermore, we confirm that permanent idiosyncratic risk still does not predict meaningful Sharpe ratios. Using the permanent idiosyncratic risk series estimated by Bayer et al. (2019) from the Panel Study of Income Dynamics (PSID), we compute a quarterly permanent risk wedge over our extended sample, which spans the financial crisis. For an EIS of 0.2, the model generates a Sharpe ratio of 0.095. Over a shorter sample, consistent with the exercise in Lettau (2002), we recover a Sharpe ratio of 0.046, close to the original estimate. These results underscore the importance of identifying the relevant group of households (the HYLW group) for explaining observed asset prices.

While our paper focuses on the time series properties of the market return, we show that

[2] Elkamhi and Jo (2023) show that *assetholders'* consumption is procyclical with respect to a variety of asset classes.

the $\beta-$wedge for HYLW households also correlates with factors pricing the cross-section of equities. Specifically, the wedge is strongly negatively correlated with the Price of Volatile Stocks (PVS; Pflueger, Siriwardane, and Sunderam, 2020), defined as the book-to-market ratio of low volatility stocks minus that of high volatility stocks. Periods of deteriorated risk sharing among these households coincide with higher required returns both for the market portfolio and across the cross-section of equities. In contrast, the $\beta-$wedge corresponding to HYHW households is largely uncorrelated with the PVS.

Finally, we extend the model to simultaneously match moments of the risk-free rate (Weil, 1989) as well as equities. We refine the baseline specification for the β-wedge to allow for offsetting movements in its mean and variance, in the spirit of prominent asset-pricing frameworks (Campbell and Cochrane, 1999; Bansal and Yaron, 2004), recently highlighted in Hassan, Mertens, and Wang (2024). The model is calibrated to match the moments of consumption growth and the $\beta-$wedge, as well as the mean and variance of the risk-free rate. For our baseline EIS of 0.2, the model-implied Sharpe ratio is 0.47, comparable to the estimate from the simpler specification. For lower values of EIS (Vissing-Jørgensen, 2002; Best, Cloyne, Ilzetzki, and Kleven, 2020), the model generates Sharpe ratios above 0.5 while continuing to match other moments.

Related Literature. A large literature examines whether heterogeneity and imperfect risk sharing can account for asset-pricing moments; see Panageas (2020) for a survey. Grossman and Shiller (1982), Mankiw (1986), Weil (1992), Constantinides and Duffie (1996), Heaton and Lucas (1996), Guvenen (2009), Kocherlakota and Pistaferri (2009), Krueger and Lustig (2010), Krusell, Mukoyama, and Smith (2011), Gârleanu and Panageas (2015), Kogan, Papanikolaou, and Stoffman (2020), Di Tella, Hébert, and Kurlat (2024), Ding and Jiang (2025), and Gomez (2025) among others, highlight conditions under which heterogeneity can meaningfully affect equilibrium risk premia.

We contribute to a long and active literature that uses consumption expenditure survey data to study asset prices (Brav, Constantinides, and Geczy, 2002; Cogley, 2002; Vissing-Jørgensen, 2002; Kocherlakota and Pistaferri, 2009), and, more recently, Elkamhi and Jo (2023). Constantinides (2025) uses CEX data to measure the welfare costs of idiosyncratic risk. An alternative approach identifies countercyclical earnings tail risk as a driver of asset prices using administrative data (see e.g., Schmidt, 2025). Green, Kogan, Papanikolaou, and Schmidt (2025) find that high-income workers at incumbent firms face large downside earnings risks from innovation at competing firms. Adams (2025) finds that high-income wage earners draw on their stock portfolios when hit with adverse non-financial income shocks, interpreting stocks as risky for this group of households. Complementary to both approaches, we isolate a subset of high-income agents whose consumption dynamics show potential for generating sizable equity premia due to uninsurable idiosyncratic risk.

The remainder of the paper is structured as follows. Section 2 presents a general heterogeneous-agent framework and derives our preferred measure of idiosyncratic risk. Section 3 investigates the mechanisms driving the large implied Sharpe ratios. Section 4 provides a two-agent framework rationalizing how both HYHW and HYLW households price assets. Section 5 concludes.

2 Theoretical Framework

We consider a consumption-based capital asset pricing model (C-CAPM) populated by a continuum of heterogeneous agents, following Nakajima (2005) and Berger, Bocola, and Dovis (2023), henceforth BBD. Time is discrete and infinite. Each agent is exposed to aggregate risk as well as uninsurable idiosyncratic risk. We denote aggregate states by ε_t and idiosyncratic states by η_t, with histories given by $\varepsilon^t = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_t)$ and $\eta^t = (\eta_0, \eta_1, \ldots, \eta_t)$. We define $\omega^t = (\eta^t, \varepsilon^t)$ as the joint history of idiosyncratic and aggregate states.

A household with joint history ω^t maximizes expected lifetime utility in the time-separable CRRA form described by equation (1), where $C(\omega^t)$ denotes the consumption, $\beta \in (0, 1)$

is the subjective discount factor, and $\gamma > 0$ is the inverse of the elasticity of intertemporal substitution (EIS).

$$\max_{\{C(\omega^t)\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{\left(C(\omega^t)\right)^{1-\gamma} - 1}{1-\gamma} \tag{1}$$

The budget constraint faced by the household is described by equation (2), where $I(\omega^t)$ denotes the income at time t, $R^f(\varepsilon^{t-1})$ is the gross return on the risk-free bond known at time $t-1$, $B(\omega^{t-1})$ are the bond holdings from the previous period, $X^i(\omega^{t-1})$ are the holdings of risky assets $i \in \mathcal{I}$ chosen in $t-1$, and $R^i(\varepsilon^t)$ denotes the gross return of asset i at time t.

$$C(\omega^t) - I(\omega^t) \leq R^f(\varepsilon^{t-1})B(\omega^{t-1}) - B(\omega^t) + \int_{\mathcal{I}} R^i(\varepsilon^t)X^i(\omega^{t-1})\,di - \int_{\mathcal{I}} X^i(\omega^t)\,di \tag{2}$$

We assume that the household is subject to the following generalized borrowing constraint:

$$\mathcal{H}\big(B(\omega^t), \{X^i(\omega^t)\}_{i\in\mathcal{I}}\big) \geq 0, \tag{3}$$

where $\mathcal{H}$ is a vector-valued function. Following BBD, a key assumption we make is that holding more risk-free bonds relaxes the borrowing constraint, thereby making it easier to borrow.

A household's stochastic discount factor (SDF), $M(\omega^{t+1})$, is defined in terms of individual consumption growth $\frac{C(\omega^{t+1})}{C(\omega^t)}$ which, however, is related to aggregate consumption growth according to:

$$\frac{C(\omega^{t+1})}{C(\omega^t)} = \frac{\delta(\omega^{t+1})}{\delta(\omega^t)} \frac{C(\varepsilon^{t+1})}{C(\varepsilon^t)}, \tag{4}$$

where $\delta(\omega^t)$ is an idiosyncratic adjustment factor that satisfies the law of large numbers:

$$\int_{\eta^t} \delta(\eta^t, \varepsilon^t)\,d\eta^t = 1 \qquad \forall\varepsilon^t,\, \forall t. \tag{5}$$

For illustration, consider first the case absent borrowing constraints. The household Euler condition is given by:

$$\mathbb{E}\big[M\big(\omega^{t+1}\big) \mid \varepsilon^t\big] = \frac{1}{R^{\mathrm{f}}(\varepsilon^t)} \tag{6}$$

where $\mathbb{E}[X(\omega^{t+1}) \mid \varepsilon^t] = \sum_{\omega^{t+1}} \Pr(\omega^{t+1} \mid \omega^t)\, X(\omega^t, \omega_{t+1})$ and $\Pr(\cdot)$ denotes transition probabilities. Substituting for the SDF, the household Euler for risk-free bonds is given by:

$$\mathbb{E}\left[\underbrace{\beta\mathbb{E}\left[\left(\frac{\delta(\eta^{t+1},\varepsilon^{t+1})}{\delta(\eta^t,\varepsilon^t)}\right)^{-\gamma}\middle|\eta^t,\varepsilon^{t+1}\right]}_{\equiv\beta(\varepsilon^{t+1},\eta^t)} \times \left(\frac{C(\varepsilon^{t+1})}{C(\varepsilon^t)}\right)^{-\gamma}\middle|\varepsilon^t\right] = \frac{1}{R^{\mathrm{f}}(\varepsilon^t)} \tag{7}$$

where $\mathbb{E}[X(\omega^{t+1}) \mid \eta^t, \varepsilon^{t+1}] = \sum_{\eta^{t+1}} \Pr(\eta^{t+1} \mid \eta^t, \varepsilon^{t+1})\, X(\eta^{t+1}, \varepsilon^{t+1})$ and $\mathbb{E}[X(\eta^t, \varepsilon^{t+1}) \mid \varepsilon^t] = \sum_{\varepsilon^{t+1}} \Pr(\varepsilon^{t+1} \mid \varepsilon^t)\, X(\eta^t, \varepsilon^{t+1})$. Intuitively, all agents trading the bond must agree on its price. Accordingly, the SDF of the "pricer," sometimes referred to as the "as-if" representative agent (e.g. Werning, 2015) is given by $\beta(\eta^t, \varepsilon^{t+1}) \times \left(\frac{C(\varepsilon^{t+1})}{C(\varepsilon^t)}\right)^{-\gamma}$.

Borrowing Constraints. One of the key takeaways of our paper is that borrowing constraints are empirically relevant for explaining the relatively high Sharpe ratio on equities in the data (see Section 3). As such, we assume agents face constraint (3) and we focus on households which are likely unconstrained in the current period (and thus participate in bond and equity markets) but who face a positive probability of becoming constrained in the future (making them behave more risk averse). We assume there is at least one agent who is unconstrained and this will be the most "patient" consumer who maximally values the risk-free bond. In the presence of borrowing constraints, the Euler condition for bonds is given by:[3]

[3]Given the assumption that $d\mathcal{H}/dB > 0$, this implies that the Euler holds with an inequality: $\mathbb{E}_t[(C(\omega^{t+1})/C(\omega^t))^{-\gamma}] \leq 1/R^f(\varepsilon^t)$, where $C(\omega^t)$ is consumption of a household with joint history $\omega^t = (\eta^t, \epsilon^t)$, and is only satisfied with an equality for the most patient household. The proof follows Proposition 1 in Nakajima (2005). However, even in a setting without assumptions regarding equation (3), Krusell et al. (2011) derive conditions under which it is the most patient, specifically the individual with the most to lose, who prices the assets.

$$\max_{\eta^t}\left\{\mathbb{E}\left[\beta(\eta^t,\varepsilon^{t+1})\times\left(\frac{C(\varepsilon^{t+1})}{C(\varepsilon^t)}\right)^{-\gamma}\middle|\varepsilon^t\right]\right\}=\frac{1}{R^f(\varepsilon^t)} \tag{8}$$

The SDF delivering the highest price for a risk-free bond depends on both on the mean of the household-specific $\beta-$wedge and, critically, the cyclicality of the $\beta-$wedge. A more countercyclical wedge implies a higher level of patience for the "as-if" representative investor.

Denote $\hat{\eta}^t$ as the argument that maximizes (8) and the corresponding maximal wedge $\hat{\beta}(\epsilon^{t+1})$ and SDF $\widehat{M}(\epsilon^{t+1})$. In this paper we ask: if the "pricers" (the bondholders) also price equity, how large is the predicted Sharpe ratio? As such, we assume the Euler for equities, is given as follows:

$$\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})R^{eq}(\varepsilon^{t+1})\middle|\varepsilon^t\right]=1 \tag{9}$$

2.1 C-CAPM representation

Having established the relevant pricing kernel above, we derive the C-CAPM representation of the model for the pricers' SDF. The log pricing kernel and the log $\beta-$wedge are given as follows:

$$\widehat{m}\left(\varepsilon^{t+1}\right)=\log\hat{\beta}\left(\varepsilon^{t+1}\right)-\gamma\Delta c(\varepsilon^{t+1}) \tag{10}$$

$$\log\hat{\beta}\left(\varepsilon^{t+1}\right)=\log\bar{\beta}+\phi\Delta c(\varepsilon^{t+1}) \tag{11}$$

where $\log\bar{\beta}$ is a constant and we specify consumption growth as an AR(1) process:

$$\Delta c_{t+1}=\mu_c+\alpha\Delta c_t+\varepsilon_{t+1} \tag{12}$$

The parameter α captures the persistence of consumption growth and $\varepsilon_{t+1}\sim iid\,\mathcal{N}\,(0,\,\sigma_\varepsilon^2)$.[4]

Equation (10) displays the log of the marginal investor's SDF and (11) assumes the

[4]In Section 3.6, we allow for a conditionally heteroskedastic process for $\beta-$wedge with an additional shock whose variance is related to the aggregate shock. Our main results are largely unchanged.

$\beta-$wedge is an affine function of consumption growth (see also Constantinides and Duffie, 1996; Lettau, 2002). The key parameter of interest is ϕ, which governs the cyclicality of the wedge: a negative sign implies the effective patience parameter induced by the $\beta-$wedge is countercyclical with respect to aggregate shocks so the discount factor is higher in recessions (when aggregate consumption growth is low or negative) and lower in booms.

To define equity, we follow the macro-finance literature and consider the aggregate market portfolio– the riskiest portfolio in the economy (Duffie, 2001). To calculate a target Sharpe ratio, we use the Fama-French CRSP value-weighted market return from the Kenneth French Data Library (French, 2026). We use this series since it approximates the aggregate U.S. market portfolio better than the S&P 500; however, the resulting Sharpe ratios are similar. In the model, by definition, the market excess return RP^{m} is perfectly correlated with a claim to aggregate consumption growth and has equal variance $\sigma_\xi = \sigma_\epsilon$ (see Lucas, 1978).[5]

Our main exercise is to derive the Sharpe ratio implied by the bondpricers' SDF. As in Lettau (2002), the Sharpe ratio is a function of the correlation between the relevant SDF and the market return, as shown below.

$$\text{Sharpe ratio} = -\rho\left(\widehat{M}(\varepsilon^{t+1}); RP^{\mathrm{m}}(\varepsilon^{t+1})\middle|\varepsilon^t\right)\frac{\sigma\left(\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right)}{\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right]} = (\gamma-\phi)\sigma_\varepsilon \tag{13}$$

The first term (γ) in equation (13) is the inverse EIS coefficient which captures standard representative-agent dynamics. To resolve the equity risk premium puzzle, Mehra and Prescott (1985) show a very low EIS is required, corresponding to a level of risk aversion of, e.g., $\gamma > 40$. The second term (ϕ) captures the effective risk aversion correction for the pricer, induced by the presence of uninsurable idiosyncratic risk (e.g. the probability of becoming constrained), which varies with the aggregate state.[6]

In the data, the measured SDF is strongly countercyclical and highly volatile for HYLW

[5]Naturally, this is an extreme assumption adopted by the literature. See discussions in Campbell and Cochrane (1999) and Jung and Marin (2026), amongst others.

[6]The RHS of (13) is exactly equal to the Sharpe on the log SDF, see Appendix B.1 for a derivation.

households. Intuitively, these households face substantial income risk and possess limited buffer-stock wealth. In recessions, when aggregate consumption growth slows, the likelihood of hitting the borrowing constraint rises and their precautionary motive strengthens. As a result, their marginal utility in bad states rises sharply. Consequently, an extra unit of consumption lost in the equity market during a downturn is valued relatively more by these households and so they demand a high equity premium.

2.2 Measuring β-wedges in the CEX

Our empirical analysis relies on two U.S. time series: log aggregate consumption growth, Δc_t, and the $\beta-$wedge. The aggregate consumption growth series is constructed using NIPA data (Bureau of Economic Analysis, 2026). The annual $\beta-$wedge series is constructed from the U.S. Consumer Expenditure Survey (US Bureau of Labor Statistics, 2025), over the period 1992–2017. Our measurement of wedges closely follows the methodology in BBD, but we highlight aspects particularly relevant to our application below.[7]

Measuring the discount factor wedge requires evaluating a conditional expectation of marginal utility growth over idiosyncratic states for unconstrained households, but in the data we can only observe a single realization every period. Instead, as in BBD, we proxy the $\beta-$wedge with the cross-sectional average of marginal utility growth across households within a group, and interpret within-group household variation as realizations of idiosyncratic risk.

For each household h at time t, we compute the relative consumption share: $\varphi_{h,t} = \frac{C_{h,t}}{C_t}$, where $C_{h,t}$ and C_t denote household h and aggregate consumption, respectively. Individual household consumption is residualized on observable demographic characteristics, as detailed below. At each date, we sort households into high and low labor-income groups, defining the threshold as the median income. Then, within each income group, we sort into below- and above-median wealth groups. This approach balances achieving sufficient granularity against maintaining sufficient sample size within the CEX. As a result, for each of the four

[7]We refer the reader to Appendix D in Berger et al. (2023) for details on sample selection.

groups of households, our sample contains approximately 875 observations per year. We present summary statistics of the underlying micro-data for each high income group, including demographic factors as well as details on their wealth and assetholdings, in Table A.1.

The $\beta-$wedge for a group g, consisting of N_g households, is constructed as follows:

$$\hat{\beta}_{g,t+1} \equiv \frac{1}{N_g} \sum_{h=1}^{N_g} \left(\frac{\varphi_{h,t+1}}{\varphi_{h,t}} \right)^{-\gamma} \tag{14}$$

and we de-mean the log of the constructed series to mitigate measurement error.[8]

We focus on the two high-income groups because they are most likely to be unconstrained and thus on their respective risk-free bond Euler equations. Amongst these, we select the high-income, low-wealth (HYLW) group because these households are the most patient – willing to pay the highest price for a risk-free bond, according to the data. To see this, rewrite the LHS of (8) as follows:

$$\max_{\eta^t} \left\{ \mathbb{E}_t \left[\beta(\eta^t, \varepsilon^{t+1}) \times \left(\frac{C(\varepsilon^{t+1})}{C(\varepsilon^t)} \right)^{-\gamma} \right] \right\} = \max_{\eta^t} \left\{ \mathbb{E}_t \left[\beta(\eta^t, \varepsilon^{t+1}) \right] \mathbb{E}_t \left[\left(\frac{C(\varepsilon^{t+1})}{C(\varepsilon^t)} \right)^{-\gamma} \right] + \mathrm{Cov}_t \left[\beta(\eta^t, \varepsilon^{t+1}), \left(\frac{C(\varepsilon^{t+1})}{C(\varepsilon^t)} \right)^{-\gamma} \right] \right\} \tag{15}$$

where $\mathbb{E}_t$ is conditional on ε^t. When the $\beta-$wedge is countercyclical – negatively covarying with aggregate consumption growth – households in this group demand a higher price for risk-free bonds. Figure 1a confirms that the $\beta-$wedge for HYLW households is strongly counter-cyclical for $EIS < 0.3$, in contrast to HYHW households whose wedge is procyclical (see Section 3.3).

We construct $\beta-$wedges corresponding to values of EIS $\gamma^{-1} \in \{1, 0.5, 0.3, 0.2\}$. Greater risk aversion increases the curvature of the utility function and disproportionately weights

[8] If measurement error is multiplicative: $C_{h,t}^{\text{measured}} = C_{h,t}^{\text{true}} \times \exp(\zeta_{h,t})$, where $\zeta_{h,t}$ is classical measurement error that is independent across households and stationary over time, then measurement error introduces a multiplicative bias in the $\beta-$wedge $\hat{\beta}_{g,t}^{\text{measured}} = \hat{\beta}_{g,t}^{\text{true}} \times \exp(\gamma^2 \sigma_\zeta^2)$. Crucially, if the bias is constant over time, this cancels when we de-mean the wedge series or examine its covariance with aggregate variables. See also Kocherlakota and Pistaferri (2009).

households experiencing low relative consumption growth – emphasizing bad realization of the idiosyncratic state. To illustrate, consider the following simplified example. A group consists of two households whose ratio of individual-to-aggregate consumption growth rates $\varphi_{h,t+1}/\varphi_{h,t}$ are 1.5 and 0.75, i.e. one is growing relatively fast, the other relatively slow. For $\gamma = 1$, the components of the $\beta-$wedge are $1.5^{-1} \approx 0.66$ and $0.75^{-1} \approx 1.33$, which average to 1 (using (14)), so the economy (and asset prices) still behaves as if there were a representative agent. For the same consumption realizations with $\gamma = 5$, the components become $1.5^{-5} \approx 0.13$ and $0.75^{-5} \approx 4.21$: the low-growth household realizations now dominate, and the $\beta-$wedge rises to approximately 2.17. The presence of idiosyncratic risk now matters quantitatively for asset prices.

We emphasize two further aspects of the measurement procedure. First, following Vissing-Jørgensen (2002) and BBD, we construct the $\beta-$wedge using semi-annual changes in consumption, which we then aggregate to annual frequency. This choice reflects two considerations. At higher frequencies, timing mismatches between when consumption occurs and when it is recorded introduce noise, a concern known as time aggregation bias. Second, households may reclassify expenditures across categories between adjacent quarterly interviews, creating spurious volatility in measured consumption growth, termed category switching. The CEX panel structure, which follows households for approximately one year, makes annual frequency the natural horizon for computing consumption growth rates.

Second, to mitigate measurement error beyond the demeaning procedure described above, we implement several additional steps following BBD: (i) winsorizing consumption at the top and bottom 1% to remove extreme outliers from reporting errors; (ii) trimming observations where semi-annual consumption growth ratios fall outside the 0.2–5 range; (iii) residualizing consumption on demographics (age, sex, race, education, state) to remove life-cycle and compositional effects; and (iv) scaling household consumption growth by NIPA aggregate consumption growth, which is arguably measured with less error than CEX-based aggregates.

3 Results

Figure 1 already summarized our main results. The left panel (Figure 1a) reports estimates of the cyclicality of the $\beta-$wedge denoted by ϕ, which is significantly negative for EIS values of 0.3 and below. The $\beta-$wedge is substantially more volatile and countercyclical than, e.g., permanent risk measures previously used in the literature. Across all specifications in Figure 1, the estimated volatility of aggregate consumption growth is 1.1%, resulting in an aggregate innovation $\sigma_\epsilon = 0.008$, in line with postwar data for the U.S. While we do not take a stance on the correct value of the EIS, Gruber (2013) and Best et al. (2020) provide evidence supporting values as low as 0.1. Lower EIS values generate larger Sharpe ratios, but may come at the expense of fitting other moments of the data. We discuss the implications for the volatility of risk-free rates (Weil, 1992) in Section 3.6, and estimate an extension of the model that resolves this issue.

For an EIS of 0.3, the model accounts for more than 40% of the observed equity premium while with an IES of 0.2, corresponding to our benchmark, the model explains almost the entirety of the Sharpe ratio. The cases $\gamma = 1$ (log utility) and $\gamma^{-1} = 0.5$ yield acyclical estimates of ϕ and imply Sharpe ratios below 0.05, accounting for less than 10% of the historical U.S. value of 0.5.[9]

3.1 Composition vs. Within-Group Variation

Based on equation (14), the $\beta-$wedge for group g can be decomposed into two components:

$$\log \hat{\beta}_{g,t+1} = \underbrace{-\gamma\left(\Delta \log C_{g,t+1} - \Delta \log C_{t+1}\right)}_{\text{Group-Specific component}} + \underbrace{\log\left(\frac{1}{N_g}\sum_{h=1}^{N_g}\left[\frac{C_{h,t+1}}{C_{h,t}}\frac{C_{g,t}}{C_{g,t+1}}\right]^{-\gamma}\right)}_{\text{Jensen's term}}, \tag{16}$$

where the first-term captures movements in the group's average consumption relative to aggregate consumption, and the second – a Jensen's term – captures dispersion in consumption

[9] See Appendix A.1 for details on the calculation of the market U.S. Sharpe ratio used in this study.

Table 1. Robustness and Mechanisms

Row		ϕ	σ_ε	Implied Sharpe Ratio	Sharpe Explained
Wedge Decomposition:					
(1)	Within-Group Jensen	-55.547	0.008	0.512	102.5%
		(21.963)			
(2)	Between-Group term	0.159	0.008	0.041	8.2%
		(0.585)			
Business-Cycle Split Exercise:					
(3)	Consumption Boom	-50.424	0.008	0.435	87.0%
		(33.640)			
(4)	Consumption Bust	-66.688	0.009	0.624	124.9%
		(39.477)			
Alternative Wedges:					
(5)	High-Income/High-Wealth group	34.952	0.008	-0.254	—
		(22.677)			
(6)	Permanent Risk wedge	-4.681	0.010	0.095	19.0%
		(1.725)			

Notes. The Table reports estimates of the countercyclicality parameter ϕ and implied Sharpe ratios across various specifications. Standard errors (in parentheses) are derived from a parametric AR(1) bootstrap procedure, detailed in Appendix A.4. The Sharpe ratio benchmark is 0.50, and all results correspond to an EIS of 0.2. Unless otherwise stated, estimates use the 1992–2017 sample. The permanent risk wedge follows Bayer et al. (2019) (1983–2013 period). The wedge decomposition is based on equation (16). The business-cycle split fits an AR(1) process to aggregate consumption growth, partitioning the sample by whether the residual is above (boom) or below (bust) the median.

growth across households, indexed by h, within the group.

Rows (1)-(2) of Table 1 report estimates using each component separately. The countercyclicality of the HYLW $\beta-$wedge is driven entirely by within-group variation; the group-specific component is acyclical. Asset prices are therefore driven by the subset of HYLW households who, in each period, face high consumption risk, rather than by the group's average consumption dynamics. The Jensen's term alone generates a Sharpe ratio of 0.51, comparable to our baseline estimate, confirming that idiosyncratic risk within this group is the operative mechanism.

3.2 Booms vs. Recessions

Relatedly, we examine whether the baseline result is driven by recessions when HYLW households are most likely to face binding borrowing constraints. We split the sample based on whether aggregate shocks —residuals from an AR(1) consumption growth process— are above ("consumption boom") or below ("consumption bust") their median, in the spirit of distinguishing good and bad aggregate states as, e.g., in Krusell and Smith (1998).

Rows (3)-(4) of Table 1 report the results for these subsamples. The estimates for ϕ remain negative in both regimes (-50.4 in booms and -66.7 in busts), indicating that the countercyclicality of the $\beta-$wedge is robust to the state of the business cycle. However, the economic magnitudes differ substantially: the model-implied Sharpe ratio is nearly 50% higher in busts than in booms (0.62 versus 0.44). It is well understood in the literature that equity returns are higher during periods of crisis (see, e.g., Martin, 2017, who measures equity premia using options data). Moreover, our finding is consistent with the theoretical mechanism we describe: high model-implied Sharpe ratios arise during periods of aggregate economic downturn when a larger share of HYLW households is likely to be constrained, and thus these households demand a higher required return for holding risky assets (see also Adams, 2025).

3.3 High-income/high-wealth group

A natural question is whether high-income, high-wealth (HYHW) households – who are likewise unconstrained and hold assets – can also rationalize observed equity premia. We construct the $\beta-$wedge for this group and find that it is procyclical, implying a *negative* Sharpe ratio across specifications. For our baseline EIS of $\gamma^{-1} = 0.2$, the estimated ϕ is 34.9 and the implied Sharpe ratio is -0.25 (Table 1, row (5)). We find negative Sharpe ratios for HYHW households for all values of EIS we consider.

Our findings suggest HYHW households are less exposed to market risk and/or rarely pushed toward their borrowing constraints, such that their effective marginal utility varies little –or procyclically– across aggregate states. However, since HYHW households do hold assets, we interpret our results to suggest that consumption growth for this group is a poor proxy for their valuations. Section 4 presents a simple framework to illustrate that both HYLW and HYHW households can agree on the price of equity, yet a procyclical $\beta-$wedge arises for the latter group because of a wealth preservation (bequest) motive. Since the HYLW households have, by construction, low wealth, this non-homotheticity does not affect the measurement of their $\beta-$wedge (14).

Table A.1 (in the Appendix) details asset ownership across the two groups for a representative year 2006. We define assetholders as CEX households who report positive holdings of financial securities at the end of the previous year.[10] Together, high-income individuals account for 62% of all assetholders. Two results stand out from this table. First, HYLW households do hold assets, accounting for 20% of assetholders. Second, the assetholder group is dominated by HYHW households (42%)– who we showed would demand a negative Sharpe ratio. As such, these findings advise caution when using assetholder consumption to inform asset prices.

3.4 Permanent-Risk Wedge

We contrast our findings with those implied by the permanent idiosyncratic risk model of Constantinides and Duffie (1996), which Lettau (2002) showed is unable to generate a sizable equity premium. Using the permanent income risk series estimated by Bayer, Luetticke, Pham-Dao, and Tjaden (2019) from the Panel Study of Income Dynamics (PSID), which extends the Storesletten, Telmer, and Yaron (2007) estimates to cover the Global Financial Crisis, we construct a quarterly permanent-risk wedge. This is scaled by $\frac{\gamma(\gamma+1)}{2}$, as implied by the permanent risk framework, and we report the annualized Sharpe ratio (scaling by $\sqrt{4}$).

For our preferred specification with an EIS of 0.2, the model generates a Sharpe ratio of only 0.095, or 19% of our empirical target (Table 1, row 6). This magnitude is directly comparable

[10]Due to data limitations, we follow the literature and do not distinguish between bond holders and equity holders. We discuss the definition of assetholders in Appendix A.3.

to Lettau (2002), who finds a Sharpe ratio of approximately 0.06.[11] Neither estimate comes close to the Sharpe ratio observed in the data. Permanent-risk, while theoretically consistent with precautionary savings motives, in the data is neither sufficiently volatile nor sufficiently countercyclical to match target Sharpe ratios. These results underscore the role of potentially binding borrowing constraints and the consumption risk faced specifically by HYLW households as a potential explanation for the observed asset prices.

3.5 Pricing the Cross-Section

So far, we have investigated the ability of our constructed idiosyncratic risk measure to account for the time-series properties of the market return. In this section, we show that the $\beta-$wedge also correlates with a factor that prices the cross-section of equities: the Price of Volatile Stocks (PVS) of Pflueger, Siriwardane, and Sunderam (2020).[12] PVS is defined as the book-to-market ratio of low volatility stocks minus the book-to-market ratio of high volatility stocks. When PVS is high, volatile stocks are relatively expensive (low book-to-market), indicating high risk appetite. When PVS is low, investors demand a large premium for holding volatile stocks, reflecting elevated perceived risk.

Figure 2 reports the correlation between the constructed $\beta-$wedges and annualized PVS over 1992–2017, excluding 2009 to ensure the relationship is not driven by extreme crisis-period observations.[13] If the consumption dynamics of HYLW households constitute the relevant pricing kernel, then their $\beta-$wedge should be negatively correlated with PVS: a high wedge signals elevated uninsurable idiosyncratic risk and greater demand for precautionary savings, which should coincide with high required returns on risky assets.

Panel A confirms this prediction. The HYLW $\beta-$wedge is strongly negatively correlated with PVS ($\rho = -0.45$), explaining approximately 20% of its variation: years in which risk

[11]This difference partly arises from our longer sample period which includes the Global Financial Crisis. Restricting the sample to 1983-1999, we recover a Sharpe ratio of 0.05.

[12]Pflueger et al. (2020) show that low PVS forecasts declines in real investment and is associated with high safe asset prices—patterns consistent with risk-centric theories of the business cycle.

[13]Including 2009 strengthens the results: the correlation for HYLW households increases to -0.60 with $R^2 = 0.36$. See Figure D.1 in the Appendix.

Figure 2. $\beta-$wedge and the Price of Volatile Stocks

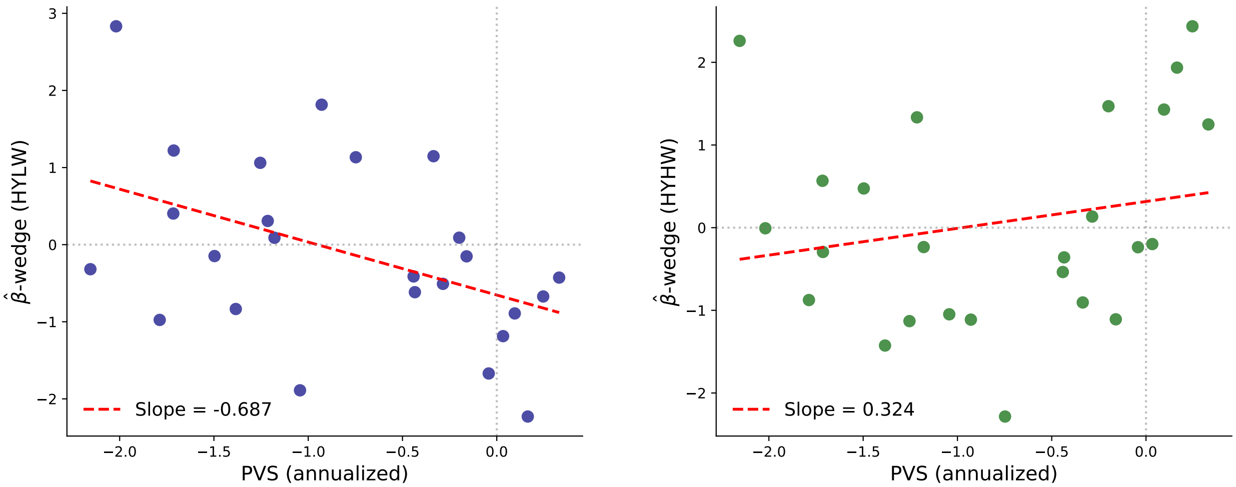

(a) HYLW **(b) HYHW**

Notes. This figure plots the $\beta-$wedge against the annualized Price of Volatile Stocks (PVS) from Pflueger et al. (2020) for 1992–2017, excluding 2009. Panel A shows HYLW (high-income, low-wealth) households; Panel B shows HYHW (high-income, high-wealth) households.

sharing deteriorates for these households coincide with years in which equity investors demand high compensation for holding volatile stocks. Panel B shows that no such relationship exists for HYHW households—their $\beta-$wedge is *positively* correlated with PVS ($\rho = 0.20$) and explains only 4% of PVS variation.

3.6 Reconciling risk-free rates

A central challenge for models reconciling the equity risk premium is to avoid introducing excessive volatility in the implied risk-free rate (Weil, 1989). In our framework, the as-if representative agent with a $\beta-$wedge, evaluated for modest values of EIS, essentially behaves as highly risk-averse investor. Under log-normality, the log risk-free rate is given by:

$$r_t^f = -\mathbb{E}[\widehat{m}_{t+1} \mid \varepsilon^t] - \frac{1}{2}\mathbb{V}\left[\widehat{m}_{t+1} \mid \varepsilon^t\right], \tag{17}$$

so a higher conditional volatility of $\hat{m}$ can lead to counterfactual predictions for moments of the risk free rate.

The consumption process is given in equation (12). We build on statistical processes used in the literature (Cox, Ingersoll, and Ross, 1985; Sun, 1992; Backus, Foresi, and Telmer, 1998), which capture features of SDFs in prevalent models (see, e.g., Campbell and Cochrane, 1999; Bansal and Yaron, 2004). Namely, we specify an extended process for the $\beta-$wedge featuring offsetting movements in its mean and variance, and a state-dependent conditional volatility:

$$\log \hat{\beta}_{t+1} = \log \bar{\beta} + \chi \sigma_{\nu,t}^2 + \phi \Delta c_{t+1} + \nu_{t+1}, \tag{18}$$

$$\nu_{t+1} \mid \varepsilon^t \sim \mathcal{N}\left(0, \sigma_{\nu,t}^2\right), \quad \sigma_{\nu,t}^2 = \mu_\sigma + \phi^\sigma \left(\Delta c_t - \frac{\mu_c}{1-\alpha}\right)., \tag{19}$$

where χ allows for the $\beta-$wedge to comove with its conditional volatility, and the parameter can be thought to capture precautionary savings, see e.g. Verdelhan (2010). We also assume $\mu_\sigma > 0$ and $\phi^\sigma < 0$ so that the conditional volatility of the $\beta-$wedge is positive on average and high during periods of low consumption growth. The Sharpe ratio continues to be given by equation (13).

Panel A of Table 2 reports our chosen parameter values. We choose μ_c, α and σ_ϵ to exactly match moments of the log consumption growth process. We choose ϕ consistent with Section 3, based on a simple OLS regression of log $\beta-$wedge on log consumption growth. We jointly choose μ_σ, ϕ^σ, and, critically, χ, to match the volatility of the log $\beta-$wedge and of the risk-free rate, and we require that $\sigma_{\nu,t}$, specified in equation (19), is negative no more than 1% of the time.[14] Finally, $\log \overline{\beta}$ is chosen to match the mean of the risk-free rate.

Panel B of Table 2 reports the implied moments. The model exactly matches the mean, variance and autocorrelation of consumption growth, as well as the mean and variance of the $\beta-$wedge. We underestimate the autocorrelation of the $\beta-$wedge, but the point estimate in the data is not significantly different from zero given the 95% confidence interval. While still

[14]Truncating the volatility process to remain positive is common practice in similar models (see, e.g., Lettau, 2002), and in any model based on a discrete time version of Cox, Ingersoll, and Ross (1985) (e.g., Backus, Foresi, and Telmer, 1998).

Table 2. Model Calibration and Unconditional Moments

Panel A: Parameters

γ^{-1}	μ_c	α	σ_ε	ϕ	χ	ϕ^σ	μ_σ	$\log\bar{\beta}$
0.2	0.008	0.672	0.008	-50.140	-1.179	-51.833	1.385	2.140

Panel B: Unconditional Moments

		EIS = 0.2	
Moment		Data	Model
$\mathbb{E}[\Delta c]$	$\mu_c(1-\alpha)^{-1}$	0.024	0.024
$\sigma[\Delta c]$	$\sigma_\varepsilon(1-\alpha^2)^{-1/2}$	0.011	0.011
$\rho_{\Delta c}$	α	0.672	0.672
$\sigma\left[\log\hat{\beta}\right]$	$\sqrt{\phi^2\sigma_\varepsilon^2+(\chi\phi^\sigma+\phi\alpha)^2\sigma_\varepsilon^2(1-\alpha^2)^{-1}+\mu_\sigma}$	1.290	1.290
$\rho_{\log\hat{\beta}}$	$\dfrac{\sigma_\varepsilon^2(\chi\phi^\sigma+\phi\alpha)(\chi\phi^\sigma\alpha+\phi)}{(1-\alpha^2)(\phi^2\sigma_\varepsilon^2+\mu_\sigma)+\sigma_\varepsilon^2(\chi\phi^\sigma+\phi\alpha)^2}$	0.374	-0.019
$\mathbb{E}[R^f]$	$\exp\left(\mu_r+\frac{1}{2}\sigma_r^2\right)-1$	0.025	0.025
$\sigma(R^f)$	$\sqrt{\exp(2\mu_r+\sigma_r^2)\left[\exp(\sigma_r^2)-1\right]}$	0.022	0.022
Sharpe Ratio	$(\gamma-\phi)\,\sigma_\varepsilon$	0.500	0.468

Panel C: Additional Expressions

(1) $R^f = e^{r^f}$

(2) $\mu_r = -\log\bar{\beta}+(\gamma-\phi)\mu_c(1-\alpha)^{-1}-\frac{1}{2}(\gamma-\phi)^2\sigma_\varepsilon^2-(\chi+\frac{1}{2})\mu_\sigma$

(3) $\sigma_r^2 = \left[(\gamma-\phi)\alpha-(\chi+\frac{1}{2})\phi^\sigma\right]^2\dfrac{\sigma_\varepsilon^2}{1-\alpha^2}$

delivering a Sharpe ratio of 0.47, this extended model also exactly matches the mean and volatility of the risk-free rate.[15]

4 HYHW Households and Wealth Preservation

Our results show that HYLW households' consumption growth is sufficiently volatile and risky to explain the observed Sharpe ratio on equity, while HYHW household consumption dynamics imply a smaller Sharpe ratios. This is a striking result because, ceteris paribus, in equilibrium HYHW households are then the *only* agents that should be holding equity. Our findings are suggestive of a model where consumption growth is not informative for HYHW households' valuation. We rationalize this in a stylized portfolio choice framework (in the spirit of Merton, 1969), extended to allow for a bequest/wealth motive (as in, e.g., De Nardi, 2004; Straub, 2019).

Consider a two-period model, $t \in \{0, 1\}$, with two agents, $i \in \{A, B\}$. Agents are identical at $t = 0$ except that agent B has a bequest motive. Agent A maximizes (20), while agent B maximizes (21), where $c_{i,t}$ denotes consumption, and a_B denotes bequests. Both u and g are CRRA with identical curvature γ.[16]

$$U_A = u(c_{A,0}) + \beta \, \mathbb{E}_0[u(c_{A,1})] \tag{20}$$

$$U_B = u(c_{B,0}) + \beta \, \mathbb{E}_0[u(c_{B,1}) + g(a_B)] \tag{21}$$

In period 0, each agent chooses consumption, bond holdings b_i, and equity holdings θ_i subject to:

$$c_{i,0} + b_i + \theta_i = w_{i,0}, \tag{22}$$

[15] Appendix C details the calibration strategy and reports calibrations with EIS values $\gamma^{-1} \in \{0.3, 0.1\}$.
[16] Specifically, $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$, and $g(x) = \frac{x^{1-\gamma}}{1-\gamma}$.

where $w_{i,0}$ denotes initial wealth. Period-1 wealth is given by:

$$w_{i,1} \;=\; R^f b_i + R\,\theta_i \;=\; \underbrace{R^f\left(w_{i,0} - c_{i,0}\right)}_{\bar{w}_i} + \theta_i \underbrace{\left(R - R^f\right)}_{X}, \tag{23}$$

where R^f and R denote risk-free and risky returns, respectively, and X is the realized excess return.

Agent A consumes all their wealth in period 1, so that $c_{A,1} = w_{A,1}$. Agent B splits wealth between consumption and bequests according to $u'(c_{B,1}) = g'(a_{B,1})$, which under our functional form assumptions implies $c_{B,1} = w_{B,1}/2$. While the intratemporal allocation affects the marginal propensity to consume (MPC) out of wealth, it does not affect the date-0 portfolio choice problem. To see this, let $\mathcal{H}(w_{B,1}) = \max\limits_{c_{B,1}+a_B=w_{B,1}} \{u(c_{B,1}) + g(a_B)\}$ denote indirect utility for agent B. The risky asset Euler equation for agent B is given below:

$$0 = \mathbb{E}_0[\mathcal{H}'(w_{B,1})X] \iff 0 = \mathbb{E}_0\left[w_{B,1}^{-\gamma} X\right], \tag{24}$$

where the second equality follows from the envelope condition $\mathcal{H}'(w) = u'(w/2)$.

We solve the portfolio choice problem using a second-order approximation. Let $\mu \equiv \mathbb{E}_0[X]$ denote the mean excess return to equity and $\sigma_X^2 \equiv \mathbb{V}_0(X)$ its variance.

Lemma 1 (Portfolio choice and savings)**.** *Both agents $i = \{A, B\}$ invest the same share of their wealth in risky assets $\theta_i \approx \frac{\bar{w}_i}{\gamma}\frac{\mu}{\sigma_X^2}$. Their savings to consumption ratios ζ_i are given by:*

$$\zeta_A \approx \left(\beta(R^f)^{1-\gamma}\left[1 - \frac{\gamma-1}{2\gamma}\frac{\mu^2}{\sigma_X^2}\right]\right)^{\frac{1}{\gamma}},$$

$$\zeta_B \approx \left(2^\gamma\beta(R^f)^{1-\gamma}\left[1 - \frac{\gamma-1}{2\gamma}\frac{\mu^2}{\sigma_X^2}\right]\right)^{\frac{1}{\gamma}},$$

respectively.

Proof. See Appendix E. □

The 2^γ in the expression for ζ_B reflects the increased effective concavity in period-1 utility: since agent B consumes only half of their wealth, the marginal value of an additional unit of savings is higher. Period-0 consumption is $c_{i,0} = w_{i,0}/(1+\zeta_i)$. With CRRA utility over both consumption and wealth, assuming the *same* curvature, the bequest motive changes the MPC out of wealth in period 1, but (locally) does not change the risky demand share. While the wealth-aware agent B faces a stronger incentive to save, the marginal benefit of consumption at time 0 is increasing, preventing overwhelming savings ($w_{B,1} < 2w_{A,1}$).

Proposition 1. *The consumption of households with a bequest motive (agent B) is less sensitive to market returns, namely:* $\mathrm{cov}_0(-\gamma c_{A,1}, R) < \mathrm{cov}_0(-\gamma c_{B,1}, R) < 0$.

Proof. See Appendix E. $\square$

Intuitively, the presence of a bequest motive implies that agent B's consumption is less exposed to fluctuations in risky asset returns than that of agent A, as agent B saves a portion of realized returns. Yet, since agent B places greater weight on wealth preservation, the marginal utilities of the two agents exhibit similar exposure to risky asset returns. As such, an econometrican applying equation (13) to HYHW misattributes their valuation because their true SDF also involves a wealth preservation motive.[17]

5 Conclusion

This paper revisits the classic question of whether idiosyncratic risk can account for the observed equity premium, bringing to bear a new measure of imperfect risk sharing grounded in recent methodological advances. Using discount-factor wedges measured directly from household-level consumption data, we find that uninsurable idiosyncratic risk among high-income, low-wealth (HYLW) households can explain 94% of the historical U.S. Sharpe ratio for an EIS of 0.2. This mechanism is driven by the countercyclicality of the "as-if" representative

[17] More generally, in incomplete-markets models with trading frictions (Di Tella et al., 2024; Kocherlakota, 2025), wealthy agents in the accumulation phase absorb aggregate shocks into their asset holdings rather than consumption, dampening the covariance between their consumption growth and market returns.

agent's marginal utility growth, as a result of idiosyncratic consumption risk of households facing a positive probability of binding borrowing constraints, particularly during downturns.

Our analysis shows that within-group household-level heterogeneity, capturing realizations of the idiosyncratic risk, is the dominant driver of the results. In contrast i) permanent-income-risk measures, even updated with post-crisis micro data from the PSID, account for only a small fraction of the observed equity premium, and ii) HYHW households face procyclical consumption risk. We rationalize the latter within a simple framework in which wealthy households' valuations are driven by a wealth preservation motive, making their consumption growth a poor proxy.

References

Adams, P. (2025). Stocks for the long run or liquidity? tax data evidence and portfolio choice implications. Working Paper.

Backus, D., S. Foresi, and C. I. Telmer (1998, September). Discrete-time models of bond pricing. NBER Working Paper No. 6736.

Bansal, R. and B. N. Lehmann (1997). Growth-optimal portfolio restrictions on asset pricing models. *Macroeconomic Dynamics 1*(2), 333–354.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance 59*(4), 1481–1509.

Bayer, C., R. Luetticke, L. Pham-Dao, and V. Tjaden (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica 87*(1), 255–290.

Berger, D., L. Bocola, and A. Dovis (2023). Imperfect risk sharing and the business cycle. *Quarterly Journal of Economics 138*(3), 1765–1815.

Best, M. C., J. S. Cloyne, E. Ilzetzki, and H. J. Kleven (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *Review of Economic Studies 87*(2), 656–690.

Brav, A., G. M. Constantinides, and C. C. Geczy (2002). Asset pricing with heterogeneous consumers and limited participation: Empirical evidence. *Journal of Political Economy 110*(4), 793–824.

Bureau of Economic Analysis (2026). National income and product accounts (nipa). U.S. Department of Commerce. Available at https://www.bea.gov/. Accessed: 2026-01-01.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), 205–251.

Cogley, T. (2002). Idiosyncratic risk and the equity premium: Evidence from the consumer expenditure survey. *Journal of Monetary Economics 49*(2), 309–334.

Constantinides, G. M. (2025). Welfare costs of idiosyncratic and aggregate consumption shocks. *Review of Asset Pricing Studies 15*(2), 103–120.

Constantinides, G. M. and D. Duffie (1996). Asset pricing with heterogeneous consumers. *Journal of Political Economy 104*(2), 219–240.

Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985). A theory of the term structure of interest rates. *Econometrica 53*(2), 385–407.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *Review of Economic Studies 71*(3), 743–768.

Di Tella, S., B. M. Hébert, and P. Kurlat (2024). Aggregation, Liquidity, and Asset Prices with Incomplete Markets. NBER Working Paper No. 32268.

Ding, X. and Z. Jiang (2025). When does idiosyncratic risk generate asset bubbles? Working paper, available at SSRN: https://ssrn.com/abstract=5262505.

Duffie, D. (2001). *Dynamic Asset Pricing Theory* (Third ed.). Princeton, NJ: Princeton University Press.

Elkamhi, R. and C. Jo (2023). Asset holders' consumption risk and tests of conditional CCAPM. *Journal of Financial Economics 148*(3), 220–244.

Fama, E. F. and K. R. French (1993, February). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

French, K. R. (2026). U.s. research returns data. https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. Accessed: 2026-01-01.

Gârleanu, N. and S. Panageas (2015). Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing. *Journal of Political Economy 123*(3), 670–685.

Gomez, M. (2025). Wealth inequality and asset prices. *Review of Economic Studies 92*(6), 3924–3967.

Green, B., L. Kogan, D. Papanikolaou, and L. Schmidt (2025). Winners and losers: Competition, creative destruction, and labor income risk. Working paper.

Grossman, S. J. and R. J. Shiller (1982). Consumption correlatedness and risk measurement in economies with non-traded assets and heterogeneous information. *Journal of Financial Economics 10*(2), 195–210.

Gruber, J. (2013). A tax-based estimate of the elasticity of intertemporal substitution. *Quarterly Journal of Finance 3*(01), 1350001.

Guvenen, F. (2009). A Parsimonious Macroeconomic Model for Asset Pricing. *Econometrica 77*(6), 1711–1750.

Hansen, L. P. and R. Jagannathan (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy 99*(2), 225–262.

Hassan, T. A., T. Mertens, and J. Wang (2024). A currency premium puzzle. Federal Reserve Bank of San Francisco Working Paper 2024-32.

Heaton, J. and D. J. Lucas (1996). Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy 104*(3), 443–487.

Jung, J. and E. A. Marin (2026). Looking for risk: Asset pricing bounds and the wealth share. Working paper.

Kocherlakota, N. and L. Pistaferri (2009). Asset Pricing Implications of Pareto Optimality with Private Information. *Journal of Political Economy 117*(3), 555–590.

Kocherlakota, N. R. (2025). Public debt bubbles, liquidity, and risk: Policy assessments based on the zero-beta interest rate. Technical report, National Bureau of Economic Research.

Kogan, L., D. Papanikolaou, and N. Stoffman (2020). Left behind: Creative destruction, inequality, and the stock market. *Journal of Political Economy 128*(3), 855–906.

Krueger, D. and H. Lustig (2010). When is Market Incompleteness Irrelevant for the Price of Aggregate Risk (and When is it Not)? *Journal of Economic Theory 145*(1), 1–41.

Krusell, P., T. Mukoyama, and A. A. Smith, Jr (2011). Asset prices in a huggett economy. *Journal of Economic Theory 146*(3), 812–844.

Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy 106*(5), 867–896.

Lettau, M. (2002). Idiosyncratic risk and volatility bounds, or can models with idiosyncratic risk solve the equity premium puzzle? *Review of Economics and Statistics 84*(2), 376–380.

Lucas, R. E. J. (1978). Asset prices in an exchange economy. *Econometrica 46*(6), 1429–1445.

Mankiw, N. G. (1986). The Equity Premium and the Concentration of Aggregate Shocks. *Journal of Financial Economics 17*(1), 211–219.

Marin, E. A. and S. R. Singh (2025). Incomplete markets and exchange rates. Working Paper 2025-11, Federal Reserve Bank of San Francisco.

Martin, I. (2017, 02). What is the expected return on the market? *Quarterly Journal of Economics 132*(1), 367–433.

Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics 15*(2), 145–161.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics 51*(3), 247–257.

Nakajima, T. (2005). A business cycle model with variable capacity utilization and demand disturbances. *European Economic Review 49*(5), 1331–1360.

Panageas, S. (2020). The Implications of Heterogeneity and Inequality for Asset Pricing. *Foundations and Trends in Finance 12*(3), 199–275.

Pflueger, C., E. Siriwardane, and A. Sunderam (2020). Financial market risk perceptions and the macroeconomy. *Quarterly Journal of Economics 135*(3), 1443–1502.

Schmidt, L. D. W. (2025). Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk. *Journal of Financial Economics 172*, 104131.

Storesletten, K., C. I. Telmer, and A. Yaron (2007). Asset pricing with idiosyncratic risk and overlapping generations. *Review of Economic Dynamics 10*(4), 519–548.

Straub, L. (2019). Consumption, savings, and the distribution of permanent income. Unpublished manuscript.

Sun, T.-s. (1992). Real and nominal interest rates: A discrete-time model and its continuous-time limit. *Review of Financial Studies 5*(4), 581–611.

US Bureau of Labor Statistics (2025). Consumer expenditure surveys (ce). https://wwwblsgov/cex/. Accessed: 2025-09-26.

Verdelhan, A. (2010). A habit-based explanation of the exchange rate risk premium. *The Journal of Finance 65*(1), 123–146.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy 110*(4), 825–853.

Weil, P. (1989). The Equity Premium Puzzle and the Risk-Free Rate Puzzle. *Journal of Monetary Economics 24*(3), 401–421.

Weil, P. (1992). Equilibrium Asset Prices with Undiversifiable Labor Income Risk. *Journal of Economic Dynamics and Control 16*(3-4), 769–790.

Werning, I. (2015). Incomplete markets and aggregate demand. NBER Working Paper No. 21432.

Appendix

A Data Appendix

A.1 Sharpe Ratio for the U.S. Market

As a benchmark, we compute the Sharpe ratio for the U.S. equity market using the annual Fama-French three-factor data from Kenneth French's data library (French, 2026) for the period 1992–2017. The Sharpe ratio is calculated as the average annual equity premium (market return minus the risk-free rate) divided by the standard deviation of excess returns:

$$\text{Sharpe Ratio (U.S.)} = \frac{T^{-1}\sum_{t=1992}^{2017}(R_t - R_t^f)}{\sigma(R_t - R_t^f)}, \tag{A.1}$$

where R_t denotes the annual market return, R_t^f is the risk-free rate, and $T = 26$ represents the sample size. For details on the construction of the market excess return series, see Fama and French (1993). This yields a Sharpe ratio of approximately 0.50, which serves as the benchmark for the analysis.

A.2 Annual Aggregate Consumption Data

To construct the annual aggregate consumption series, we follow the methodology of Berger, Bocola, and Dovis (2023). We obtain nominal Personal Consumption Expenditures (PCE) for Nondurables and Services (from Table 2.3.5) from the Bureau of Economic Analysis (BEA) National Income and Product Accounts (NIPA) at a quarterly frequency (Bureau of Economic Analysis, 2026). We first adjust these quarterly nominal series to real terms using the corresponding PCE price indexes (from Table 2.3.4). Next, we aggregate the quarterly data to an annual frequency by calculating the simple arithmetic mean of the four quarters within each calendar year.

A.3 Summary Statistics for CEX households

Table A.1 provides some summary statistics of our micro-data for the year 2006, detailing the full sample and the two high-income groups.

Table A.1. Average Characteristics of Households in 2006 by Household Income-Wealth Classification

	Full Sample	HYLW	HYHW
Age of head	44.32	46.48	44.00
Household size	2.75	2.28	2.01
Head with college (%)	36.87	39.67	36.23
Consumption expenditures per person	10511.51	11653.26	14739.22
Labor income per person	28774.22	36274.20	45385.92
Disposable income per person	28672.91	34978.22	43296.56
Hours worked per person	1330.98	1589.24	1787.31
Wage per hour	28.80	39.61	29.26
Household's net worth	172962.68	65232.17	359914.11
Liquid assets	16573.65	6757.37	37656.66
Share of all assetholders	—	19.39%	41.54%

Notes. The sample size for the 2006 interview cohort is 2327 households. All statistics are computed using sample weights. All monetary variables are expressed in 2000 U.S. dollars. Assetholders are CEX households reporting positive financial security holdings at the end of the previous year. "Share of all assetholders" is the distribution of the total population of assetholders across these categories.

Figure A.1 plots the consumption growth for HYLW and HYHW households and the aggregate. HYHW household consumption growth correlates more strongly with aggregate consumption than that of the HYLW group (0.23 vs. 0.13).

Assetholders in the CEX Following Vissing-Jørgensen (2002) and Berger, Bocola, and Dovis (2023), we classify a household as an assetholder if it reports positive holdings of financial securities at the end of the previous year. We construct lagged securities wealth as `SECESTX` minus `COMPSECX`, where `SECESTX` denotes the estimated value of securities held at the end of the prior month and `COMPSECX` captures the reported change in that value from one year ago. Households with a positive constructed value are identified as assetholders.

Figure A.1. Standardized Annual Consumption Growth

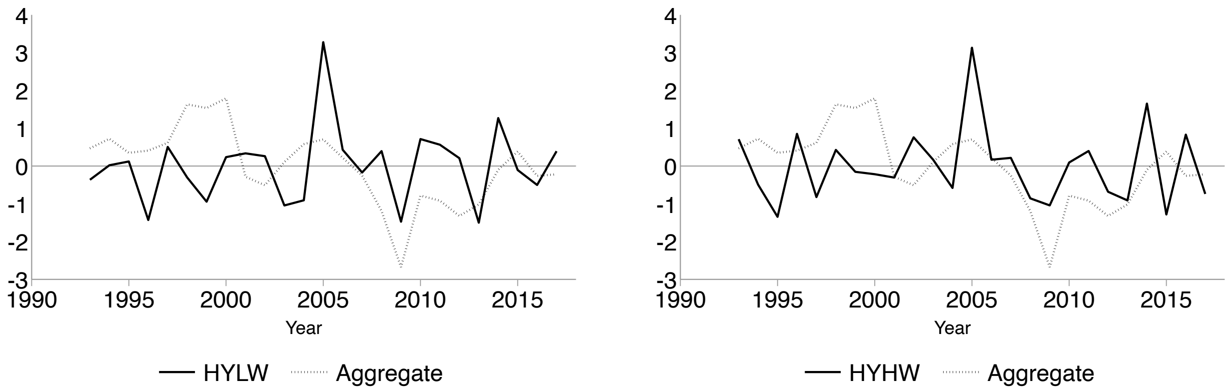

Notes. This figure reports the annual log consumption growth for selected household groups. All series are standardized. The high-income/low-net-worth and high-income/high-net-worth groups are constructed using CEX data following Berger, Bocola, and Dovis (2023). Aggregate U.S. consumption is based on NIPA data on non-durable goods and services (Bureau of Economic Analysis, 2026). The sample period is 1993–2017.

A.4 Parametric Bootstrap Procedure

To account for a possible serial correlation in our baseline estimates, given limited sample size, we employ a parametric AR(1) bootstrap procedure. First, we project the $\beta-$wedge, $\log \hat{\beta}_t$, onto consumption growth, Δc_t, and extract the ordinary least squares residuals, ν_t:

$$\log \hat{\beta}_t = \log \bar{\beta} + \phi \Delta c_t + \nu_t. \tag{A.2}$$

We then estimate an AR(1) process for these residuals:

$$\nu_t = \mu_\nu + \vartheta \nu_{t-1} + o_t, \tag{A.3}$$

to obtain the persistence parameter ϑ and the volatility of the innovations σ_o. Using these estimates, we simulate a pseudo-residual series, ν_t^b. To ensure stationarity, the initial value, ν_1^b, is drawn from the unconditional normal distribution $\mathcal{N}(0, \sigma_o^2/(1-\vartheta^2))$, while subsequent values are generated recursively using the estimated AR(1) dynamics and independent structural shocks drawn from $\mathcal{N}(0, \sigma_o^2)$. Next, we construct a bootstrapped $\beta-$wedge series, $\log \hat{\beta}^b$, by adding our simulated residuals ν_t^b to the fitted values from the initial baseline estimation:

$$\log \hat{\beta}_t^b = \widehat{\log \bar{\beta}} + \widehat{\phi}\Delta c_t + \nu_t^b. \tag{A.4}$$

We then re-estimate the baseline regression of $\log \hat{\beta}^b$ on Δc to obtain a bootstrapped coefficient, ϕ^b:

$$\log \hat{\beta}_t^b = \log \bar{\beta} + \phi^b \Delta c_t + \nu_t. \tag{A.5}$$

This procedure is repeated 2,000 times, and the standard error of our baseline estimate is computed as the standard deviation of the resulting 2,000 (B) bootstrapped coefficients $\hat{\phi}^b$:

$$SE_{boot}(\hat{\phi}) = \sqrt{\frac{1}{B-1}\sum_{b=1}^{B}\left(\hat{\phi}^b - \overline{\hat{\phi}^b}\right)^2}. \tag{A.6}$$

B Mathematical Appendix

B.1 Hansen-Jagannathan bounds and the Sharpe-ratio with idiosyncratic risk

We rederive the Sharpe ratio within our heterogeneous-agents setting. In the presence of borrowing constraints such as (3), assuming that holdings of risk-free bonds relax the constraint ($\partial\mathcal{H}/\partial B > 0$), the Euler equation for risk-free bonds holds with an inequality (Nakajima, 2005; Berger et al., 2023; Marin and Singh, 2025):

$$\mathbb{E}\left[M(\omega^{t+1})\middle|\varepsilon^t\right] \leq \frac{1}{R^f(\varepsilon^t)}. \tag{B.1}$$

The equation holds with equality for the most patient investor, characterized by (8), whose SDF we denote by $\widehat{M}(\varepsilon^{t+1})$. If the same investor prices equities, then:

$$\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})R(\varepsilon^{t+1})\middle|\varepsilon^t\right] = 1. \tag{B.2}$$

In this case, both Euler equations bind for the "pricer."

Defining the risk premium as $RP(\varepsilon^{t+1}) = R(\varepsilon^{t+1}) - R^f(\varepsilon^t)$, we obtain

$$
\begin{aligned}
\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})R(\varepsilon^{t+1})\middle|\varepsilon^t\right] &= \mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})\left(RP(\varepsilon^{t+1}) + R^f(\varepsilon^t)\right)\middle|\varepsilon^t\right] \\
&= \mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})RP(\varepsilon^{t+1})\middle|\varepsilon^t\right] + \mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})R^f(\varepsilon^t)\middle|\varepsilon^t\right] \\
&= \mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})RP(\varepsilon^{t+1})\middle|\varepsilon^t\right] + 1 \\
&= \mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right]\mathbb{E}\left[RP(\varepsilon^{t+1})|\varepsilon^t\right] \\
&\quad + \operatorname{cov}\left(\widehat{M}(\varepsilon^{t+1}); RP(\varepsilon^{t+1})\middle|\varepsilon^t\right) + 1 = 1.
\end{aligned}
$$

Moving the covariance term to the right-hand side and dividing by the product of the conditional standard deviations of the SDF and the risk premium yields:

$$
\frac{\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right]}{\sigma\left(\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right)}\frac{\mathbb{E}\left[RP(\varepsilon^{t+1})|\varepsilon^t\right]}{\sigma\left(RP(\varepsilon^{t+1})|\varepsilon^t\right)} = -\rho\left(\widehat{M}(\varepsilon^{t+1}); RP(\varepsilon^{t+1})\middle|\varepsilon^t\right) \tag{B.3}
$$

Rearranging, we obtain the Sharpe ratio as a function of the correlation between the "pricer's" SDF and the risk-premium:

$$
\underbrace{\frac{\mathbb{E}\left[RP(\varepsilon^{t+1})|\varepsilon^t\right]}{\sigma\left(RP(\varepsilon^{t+1})|\varepsilon^t\right)}}_{\text{Sharpe Ratio}} = -\rho\left(\widehat{M}(\varepsilon^{t+1}); RP(\varepsilon^{t+1})\middle|\varepsilon^t\right)\frac{\sigma\left(\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right)}{\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right]}. \tag{B.4}
$$

Combining (B.1) and (B.2) yields:

$$
\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})RP(\varepsilon^{t+1})\middle|\varepsilon_t\right] = 0. \tag{B.5}
$$

Combining (10) and (11) yields the "pricer's" SDF in the form of:

$$
\widehat{M}(\varepsilon^{t+1}) = \bar{\beta}\, e^{-(\gamma-\phi)\Delta c(\varepsilon^{t+1})}. \tag{B.6}
$$

Equation (B.7) gives the coefficient of variation implied by the lognormal structure of the SDF in equation (B.6):

$$\frac{\sigma\left[\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right]}{\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right]} = \left(e^{\mathbb{V}[-(\gamma-\phi)\Delta c(\varepsilon^{t+1})|\varepsilon^t]} - 1\right)^{1/2}. \tag{B.7}$$

Since the risky asset payoff is also lognormally distributed, the conditional correlation between the SDF component and the risky asset payoff is

$$\rho\left(\widehat{M}(\varepsilon^{t+1}),\, RP(\varepsilon^{t+1})\middle|\varepsilon^t\right) = \frac{e^{\mathrm{Cov}[-(\gamma-\phi)\varepsilon_{t+1},\xi_{t+1}]} - 1}{\left(e^{\mathbb{V}[-(\gamma-\phi)\varepsilon_{t+1}|\varepsilon^t]} - 1\right)^{1/2}\left(e^{\sigma_\xi^2} - 1\right)^{1/2}}. \tag{B.8}$$

Combining equations (B.4), (B.7), and (B.8) and taking a first-order approximation for small σ_ε^2 yields

$$\text{Sharpe ratio} = \frac{\mathbb{E}\left[RP(\varepsilon^{t+1})|\varepsilon^t\right]}{\sigma\left(RP(\varepsilon^{t+1})|\varepsilon^t\right)} \propto -\frac{e^{-(\gamma-\phi)\sigma_\varepsilon^2} - 1}{\sqrt{e^{\sigma_\varepsilon^2} - 1}} \approx -\frac{-(\gamma-\phi)\sigma_\varepsilon^2}{\sigma_\varepsilon} = (\gamma-\phi)\sigma_\varepsilon. \tag{B.9}$$

Exact Derivation for the Sharpe ratio of the log SDF. Alternatively, we can arrive at the same conclusion by deriving the log Sharpe ratio. Let lower-case letters denote the natural logarithms of gross returns, such that $r^f(\varepsilon^t) = \log(R^f(\varepsilon^t))$ and $r(\varepsilon^{t+1}) = \log(R(\varepsilon^{t+1}))$, and let the log SDF be $\widehat{m}(\varepsilon^{t+1}) = \log\widehat{M}(\varepsilon^{t+1})$. We have established that for the pricing agent, the Euler equations hold with equality: $\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})\middle|\varepsilon^t\right] = 1/R^f(\varepsilon^t)$ and $\mathbb{E}\left[\widehat{M}(\varepsilon^{t+1})R(\varepsilon^{t+1})\middle|\varepsilon^t\right] = 1$.

Applying the assumption of joint log-normality to the risk-free Euler equation,

$$\mathbb{E}\left[\exp(\widehat{m}(\varepsilon^{t+1}))\middle|\varepsilon^t\right]\exp(r^f(\varepsilon^t)) = 1, \tag{B.10}$$

and taking the natural logarithm yields:

$$r^f(\varepsilon^t) = -\mathbb{E}\left[\widehat{m}(\varepsilon^{t+1})\middle|\varepsilon^t\right] - \frac{1}{2}\mathbb{V}\left[\widehat{m}(\varepsilon^{t+1})\middle|\varepsilon^t\right]. \tag{B.11}$$

Similarly, applying the log-normal property to the risky asset Euler equation gives:

$$
\begin{aligned}
0 &= \log \mathbb{E}\left[\exp\left(\widehat{m}(\varepsilon^{t+1}) + r(\varepsilon^{t+1})\right)\middle|\varepsilon^t\right] \\
&= \mathbb{E}\left[\widehat{m}(\varepsilon^{t+1})\middle|\varepsilon^t\right] + \mathbb{E}\left[r(\varepsilon^{t+1})\middle|\varepsilon^t\right] + \frac{1}{2}\mathbb{V}\left[\widehat{m}(\varepsilon^{t+1})\middle|\varepsilon^t\right] \\
&\quad + \frac{1}{2}\mathbb{V}\left[r(\varepsilon^{t+1})\middle|\varepsilon^t\right] + \mathrm{Cov}\left(\widehat{m}(\varepsilon^{t+1}), r(\varepsilon^{t+1})\middle|\varepsilon^t\right). \qquad \text{(B.12)}
\end{aligned}
$$

Substituting Equation (B.11) into Equation (B.12) allows us to replace the expectation and variance of the SDF with $-r^f(\varepsilon^t)$:

$$
\mathbb{E}\left[r(\varepsilon^{t+1})\middle|\varepsilon^t\right] - r^f(\varepsilon^t) + \frac{1}{2}\mathbb{V}\left[r(\varepsilon^{t+1})\middle|\varepsilon^t\right] = -\,\mathrm{Cov}\left(\widehat{m}(\varepsilon^{t+1}), r(\varepsilon^{t+1})\middle|\varepsilon^t\right). \qquad \text{(B.13)}
$$

To define the Sharpe ratio, we divide both sides of Equation (B.13) by the conditional volatility of the risky asset, $\sigma\left(r(\varepsilon^{t+1})|\varepsilon^t\right)$:

$$
\begin{aligned}
&\frac{\mathbb{E}\left[r(\varepsilon^{t+1})|\varepsilon^t\right] - r^f(\varepsilon^t)}{\sigma\left(r(\varepsilon^{t+1})|\varepsilon^t\right)} + \frac{\mathbb{V}\left[r(\varepsilon^{t+1})|\varepsilon^t\right]}{2\sigma\left(r(\varepsilon^{t+1})|\varepsilon^t\right)} = -\frac{\mathrm{Cov}\left(\widehat{m}(\varepsilon^{t+1}), r(\varepsilon^{t+1})|\varepsilon^t\right)}{\sigma\left(r(\varepsilon^{t+1})|\varepsilon^t\right)} \\
&\underbrace{\frac{\mathbb{E}\left[r(\varepsilon^{t+1})|\varepsilon^t\right] - r^f(\varepsilon^t)}{\sigma\left(r(\varepsilon^{t+1})|\varepsilon^t\right)}}_{\text{Sharpe Ratio}} + \underbrace{\tfrac{1}{2}\sigma\left(r(\varepsilon^{t+1})\middle|\varepsilon^t\right)}_{\text{Jensen's Correction}} = -\rho\left(\widehat{m}(\varepsilon^{t+1}), r(\varepsilon^{t+1})\middle|\varepsilon^t\right)\sigma\left(\widehat{m}(\varepsilon^{t+1})\middle|\varepsilon^t\right). \qquad \text{(B.14)}
\end{aligned}
$$

Because the correlation coefficient ρ is bounded below by -1, the maximum achievable log Sharpe ratio is governed by the volatility of the log SDF. Using the structural assumptions on the SDF introduced above (equation (B.6)) and noting that excess risky returns are perfectly correlated with a claim to aggregate consumption (i.e., $\rho = -1$, $rp(\xi^{t+1}) = \mathbb{E}\left[rp(\xi^{t+1})|\varepsilon^t\right] + \xi_{t+1}$, and $\sigma_\xi = \sigma_\varepsilon$), the model-implied maximum Sharpe ratio simplifies exactly to:

$$
\text{Sharpe ratio} = (\gamma - \phi)\sigma_\varepsilon. \qquad \text{(B.15)}
$$

C Details for Section 3.6

For the log consumption growth rate (12), the implied unconditional mean, variance, and first-order autocorrelation are given as follows:

$$\mathbb{E}[\Delta c] = \frac{\mu_c}{1-\alpha}, \tag{C.1}$$

$$\mathbb{V}[\Delta c] = \frac{\sigma_\varepsilon^2}{1-\alpha^2}, \tag{C.2}$$

$$\rho_{\Delta c} = \alpha. \tag{C.3}$$

Equation (18) specifies the evolution of the log β-wedge. The process for the state-dependent conditional variance $\sigma_{\nu,t}^2$ is given in equation (C.4), where $\mu_\sigma > 0$ and $\phi^\sigma < 0$ imply that conditional uncertainty in the β-wedge is countercyclical. The variance $\sigma_{\nu,t}^2$ is measurable with respect to ε^t, so agents know the dispersion of β-wedge shocks before ε_{t+1} is revealed.

$$\sigma_{\nu,t}^2 = \mu_\sigma + \phi^\sigma \left(\Delta c_t - \frac{\mu_c}{1-\alpha} \right) \tag{C.4}$$

Then, the unconditional variance and first-order autocorrelation of the β-wedge are given as follows:

$$\mathbb{V}\left[\log \hat{\beta}\right] = \phi^2 \sigma_\varepsilon^2 + (\chi\phi^\sigma + \phi\alpha)^2 \frac{\sigma_\varepsilon^2}{1-\alpha^2} + \mu_\sigma, \tag{C.5}$$

$$\rho_{\log \hat{\beta}} = \frac{(\chi\phi^\sigma + \phi\alpha)(\chi\phi^\sigma\alpha + \phi)\,\sigma_\varepsilon^2}{(1-\alpha^2)(\phi^2\sigma_\varepsilon^2 + \mu_\sigma) + (\chi\phi^\sigma + \phi\alpha)^2\,\sigma_\varepsilon^2}. \tag{C.6}$$

We next derive expressions for the mean and volatility of the risk-free rate. The log stochastic discount factor is specified in equation (C.7). Under log-normality, the risk-free rate is given by equation (C.8). Substituting the log SDF into the expression for the risk-free rate yields the model-implied risk-free rate in equation (C.9). We define the gross risk-free return in levels in equation (C.10).

$$\hat{m}_{t+1} = \log \bar{\beta} + \chi \sigma^2_{\nu,t} - (\gamma - \phi)\Delta c_{t+1} + \nu_{t+1} \tag{C.7}$$

$$r^f_t = -\mathbb{E}[\hat{m}_{t+1} \mid \varepsilon^t] - \tfrac{1}{2}\mathbb{V}[\hat{m}_{t+1} \mid \varepsilon^t] \tag{C.8}$$

$$r^f_t = -\log \bar{\beta} - \chi \sigma^2_{\nu,t} + (\gamma - \phi)(\mu_c + \alpha \Delta c_t) - \tfrac{1}{2}(\gamma - \phi)^2 \sigma^2_\varepsilon - \tfrac{1}{2}\sigma^2_{\nu,t} \tag{C.9}$$

$$R^f_t \equiv \exp(r^f_t) - 1 \tag{C.10}$$

Under the specification for the conditional variance, the log risk-free rate is normally distributed with variance given in equation (C.11). The corresponding unconditional mean of the log risk-free rate is given in equation (C.12). The implied unconditional variance of the gross risk-free return is therefore given in equation (C.13).

$$\sigma^2_r = \Big[(\gamma - \phi)\alpha - (\chi + \tfrac{1}{2})\phi^\sigma\Big]^2 \frac{\sigma^2_\varepsilon}{1 - \alpha^2} \tag{C.11}$$

$$\mu_r = -\log \bar{\beta} + (\gamma - \phi)\frac{\mu_c}{1 - \alpha} - \tfrac{1}{2}(\gamma - \phi)^2 \sigma^2_\varepsilon - (\chi + \tfrac{1}{2})\mu_\sigma \tag{C.12}$$

$$\mathbb{V}(R^f_t) = \exp(2\mu_r + \sigma^2_r)\big[\exp(\sigma^2_r) - 1\big] \tag{C.13}$$

We calibrate the model in four steps. First, we calibrate $\{\mu_c, \alpha, \sigma_\varepsilon\}$ directly from consumption data moments, ensuring that the model-implied unconditional mean, standard deviation, and first-order autocorrelation of consumption growth in equations (C.1)–(C.3) match their sample counterparts exactly. Second, we set ϕ to the value obtained from the main estimation, reported in Figure 1a. Third, we select $\{\mu_\sigma, \phi^\sigma, \chi\}$ to jointly match the volatility of the β-wedge and of the risk-free rate (C.5), as well as ensuring non-negativity of the conditional variance. Because $\phi^\sigma < 0$, the conditional variance $\sigma^2_{\nu,t} = \mu_\sigma + \phi^\sigma(\Delta c_t - \frac{\mu_c}{1-\alpha})$ can turn negative for sufficiently high consumption growth realizations. We ensure σ^2_ν falls below zero no more than 1% of periods in a Monte Carlo simulation of 100,000 draws of ε from $\mathcal{N}(0, \sigma^2_\varepsilon)$. Fourth, given all remaining parameters, we recover $\log \bar{\beta}$ analytically as the unique value that sets the unconditional mean of the gross risk-free return equal to its target.

The results are reported in Table C.1.

Table C.1. Model Calibration and Unconditional Moments

Panel A: Parameters

γ^{-1}	μ_c	α	σ_ε	ϕ	χ	ϕ^σ	μ_σ	$\log\bar{\beta}$
0.3	0.008	0.672	0.008	-21.330	-1.963	-10.049	0.294	0.981
0.2	0.008	0.672	0.008	-50.140	-1.179	-51.833	1.385	2.140
0.1	0.008	0.672	0.008	-123.271	-1.625	-81.253	7.909	11.460

Panel B: Unconditional Moments

		EIS = 0.3		EIS = 0.2		EIS = 0.1	
Moment		Data	Model	Data	Model	Data	Model
$\mathbb{E}[\Delta c]$	$\mu_c(1-\alpha)^{-1}$	0.024	0.024	0.024	0.024	0.024	0.024
$\sigma[\Delta c]$	$\sigma_\varepsilon(1-\alpha^2)^{-1/2}$	0.011	0.011	0.011	0.011	0.011	0.011
$\rho_{\Delta c}$	α	0.672	0.672	0.672	0.672	0.672	0.672
$\sigma\left[\log\hat{\beta}\right]$	$\sqrt{\phi^2\sigma_\varepsilon^2+(\chi\phi^\sigma+\phi\alpha)^2\sigma_\varepsilon^2(1-\alpha^2)^{-1}+\mu_\sigma}$	0.575	0.575	1.290	1.290	3.053	3.053
$\rho_{\log\hat{\beta}}$	$\dfrac{\sigma_\varepsilon^2(\chi\phi^\sigma+\phi\alpha)(\chi\phi^\sigma\alpha+\phi)}{(1-\alpha^2)(\phi^2\sigma_\varepsilon^2+\mu_\sigma)+\sigma_\varepsilon^2(\chi\phi^\sigma+\phi\alpha)^2}$	0.313	-0.017	0.374	-0.019	0.463	-0.023
$\mathbb{E}[R^f]$	$\exp\left(\mu_r+\frac{1}{2}\sigma_r^2\right)-1$	0.025	0.025	0.025	0.025	0.025	0.025
$\sigma(R^f)$	$\sqrt{\exp(2\mu_r+\sigma_r^2)\left[\exp(\sigma_r^2)-1\right]}$	0.022	0.022	0.022	0.022	0.022	0.022
Sharpe Ratio	$(\gamma-\phi)\,\sigma_\varepsilon$	0.500	0.209	0.500	0.468	0.500	1.131

Panel C: Additional Expressions

(1) $\mu_r = -\log\bar{\beta}+(\gamma-\phi)\mu_c(1-\alpha)^{-1}-\frac{1}{2}(\gamma-\phi)^2\sigma_\varepsilon^2-(\chi+\frac{1}{2})\mu_\sigma$

(2) $\sigma_r^2 = \left[(\gamma-\phi)\alpha-(\chi+\frac{1}{2})\phi^\sigma\right]^2\dfrac{\sigma_\varepsilon^2}{1-\alpha^2}$

D Additional Results

We provide results for the full sample, including the extreme outlier in the GFC, below. Our main conclusions are unchanged.

Figure D.1. $\beta-$wedge and the Price of Volatile Stocks: Full Sample Including 2009

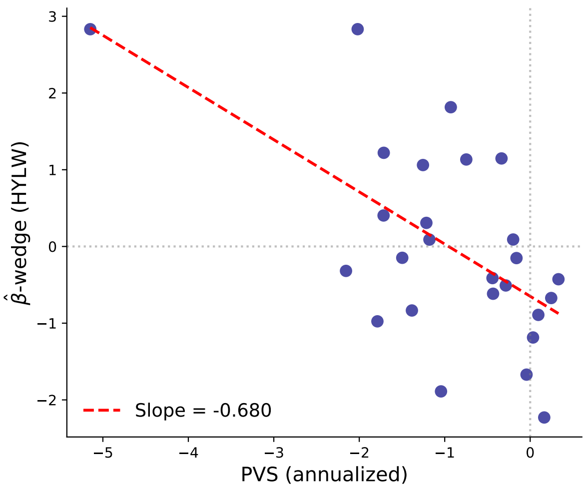

(a) HYLW

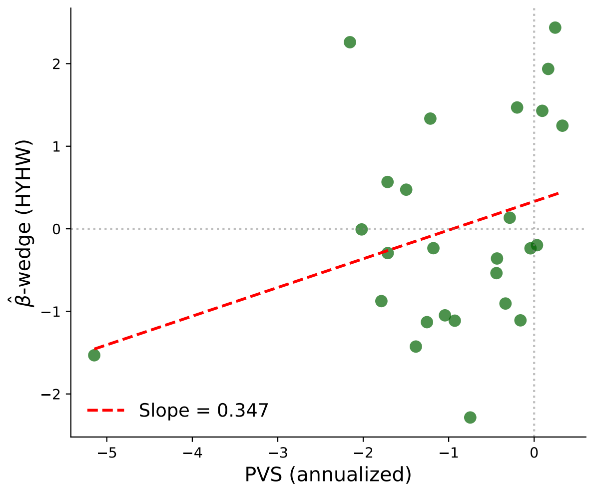

(b) HYHW

Notes. This figure replicates Figure 2 using the full sample including 2009. Panel A shows HYLW (high-income, low-wealth) households; Panel B shows HYHW (high-income, high-wealth) households. The R^2 is 0.36 and 0.10 for panels A and B respectively.

E Proofs

Proof of Lemma 1. Let $X \equiv R - R^f$ denote the excess return, with $\mu \equiv \mathbb{E}_0[X]$ and $\sigma_X^2 \equiv \mathbb{V}_0(X)$. Given time-0 wealth $\bar{w}_i \equiv R^f(w_{i,0} - c_{i,0})$, time-1 wealth can be written as:

$$w_{i,1} \;=\; \bar{w}_i + \theta_i X. \tag{E.1}$$

Agent A chooses θ_A to solve $\max_{\theta_1} \mathbb{E}_0\big[u(w_{A,1})\big]$ so the risky Euler is given by:

$$0 \;=\; \mathbb{E}_0[u'(w_{A,1})X] \;=\; \mathbb{E}_0\big[w_{A,1}^{-\gamma} X\big]\,. \tag{E.2}$$

Agent B has terminal indirect utility

$$\mathcal{H}(w_{B,1}) = \max_{c_{B,1}+a_B=w_{B,1}} \{u(c_{B,1}) + g(a_B)\}, \tag{E.3}$$

so the risky Euler is given by:

$$0 \;=\; \mathbb{E}_0[H'(w_{B,1})X]\,. \tag{E.4}$$

Under CRRA with identical curvature for u and g, the intratemporal condition implies $c = a = w/2$ and, by the envelope theorem,

$$H'(w) = u'(w/2) = (w/2)^{-\gamma} = 2^\gamma w^{-\gamma}. \tag{E.5}$$

Therefore,

$$0 \;=\; 2^\gamma \mathbb{E}_0\big[w_{B,1}^{-\gamma} X\big] \quad\Longleftrightarrow\quad 0 \;=\; \mathbb{E}_0\big[w_{B,1}^{-\gamma} X\big]\,. \tag{E.6}$$

To obtain the Merton-style approximation, expand $w_{i,1}^{-\gamma} = (\bar{w}_i + \theta_i X)^{-\gamma}$ to first order in $\theta_i X/\bar{w}_i$:

$$(\bar{w}_i + \theta_i X)^{-\gamma} = \bar{w}_i^{-\gamma}\Big(1 + \tfrac{\theta_i}{\bar{w}_i} X\Big)^{-\gamma} \approx \bar{w}_i^{-\gamma}\Big(1 - \gamma \tfrac{\theta_i}{\bar{w}_i} X\Big). \tag{E.7}$$

Substituting into $\mathbb{E}_0[w_{i,1}^{-\gamma} X] = 0$ yields

$$0 \approx \bar{w}_i^{-\gamma}\Big(\mu - \gamma \tfrac{\theta_i}{\bar{w}_i}\mathbb{E}_0[X^2]\Big), \tag{E.8}$$

so

$$\frac{\theta_i}{\bar{w}_i} \;\approx\; \frac{1}{\gamma}\frac{\mu}{\mathbb{E}_0[X^2]} \;\approx\; \frac{1}{\gamma}\frac{\mu}{\sigma_X^2}, \tag{E.9}$$

which implies the stated portfolio rule $\theta_i \approx \bar{w}_i\, \mu/(\gamma\, \sigma_X^2)$ for $i \in \{A, B\}$.

Next, we take a second-order approximation for $\mathbb{E}[w_{i,1}^{-\gamma}]$:

$$\mathbb{E}\big[w_{i,1}^{-\gamma}\big] = \mathbb{E}\big[(\bar{w}_i + \theta_i X)^{-\gamma}\big] \approx \bar{w}_i^{-\gamma}\left[1 - \gamma\frac{\theta_i}{\bar{w}_i}\mu + \frac{\gamma(\gamma+1)}{2}\frac{\theta_i^2}{\bar{w}_i^2}\mathbb{E}[X^2]\right]. \tag{E.10}$$

Substituting θ_i yields

$$\mathbb{E}_0\left[w_{i,1}^{-\gamma}\right] \approx \bar{w}_i^{-\gamma}\left[1 - \frac{\gamma-1}{2\gamma}\,\frac{\mu^2}{\mathbb{E}[X^2]}\right] \approx \bar{w}_i^{-\gamma}\left[1 - \frac{\gamma-1}{2\gamma}\,\frac{\mu^2}{\sigma_X^2}\right]. \tag{E.11}$$

Using agent A's bond Euler

$$c_{A,0}^{-\gamma} = \beta R^f\, \mathbb{E}_0\left[w_{A,1}^{-\gamma}\right], \tag{E.12}$$

substituting $\bar{w}_A = R^f(w_{A,0} - c_{A,0})$, and rearranging yields

$$\frac{w_{A,0} - c_{A,0}}{c_{A,0}} \;\approx\; \left(\beta (R^f)^{1-\gamma}\left[1 - \frac{\gamma-1}{2\gamma}\,\frac{\mu^2}{\sigma_X^2}\right]\right)^{1/\gamma} \;\equiv\; \zeta_A, \tag{E.13}$$

which delivers the result.

For agent B, the bond Euler uses $H'(w) = 2^\gamma w^{-\gamma}$. Analogous steps imply

$$\frac{w_{B,0} - c_{B,0}}{c_{B,0}} \;\approx\; \left(2^\gamma \beta (R^f)^{1-\gamma}\left[1 - \frac{\gamma-1}{2\gamma}\,\frac{\mu^2}{\sigma_X^2}\right]\right)^{1/\gamma} \;\equiv\; \zeta_B. \tag{E.14}$$

□

Proof of Proposition 1. Define

$$c_{A,1} = \overline{w}_A + \theta_A(R - R^f) \tag{E.15}$$

and

$$c_{B,1} = \frac{1}{2}[\overline{w}_B + \theta_B(R - R^f)], \tag{E.16}$$

so

$$\frac{\partial c_{A,1}}{\partial R} = \theta_A \qquad \text{and} \qquad \frac{\partial c_{B,1}}{\partial R} = \theta_B/2. \tag{E.17}$$

Plugging into the covariances yields

$$\mathrm{cov}_0(-\gamma c_{A,1}, R) = -\gamma \theta_A \mathbb{V}_0(R) \tag{E.18}$$

and

$$\mathrm{cov}_0(-\gamma c_{B,1}, R) = -\tfrac{\gamma}{2}\theta_B \mathbb{V}_0(R). \tag{E.19}$$

The result follows as long as

$$w_{B,1} < 2w_{A,1}. \tag{E.20}$$

Since

$$\bar{w}_i = R^f(w_{i,0} - c_{i,0}) = R^f \frac{\zeta_i}{1+\zeta_i} w_{i,0}, \tag{E.21}$$

and assuming

$$w_{A,0} = w_{B,0} \equiv w_0, \tag{E.22}$$

one gets

$$\frac{\bar{w}_B}{\bar{w}_A} = \frac{w_0 - c_{B,0}}{w_0 - c_{A,0}} = \frac{\frac{\zeta_B}{1+\zeta_B}}{\frac{\zeta_A}{1+\zeta_A}} \approx 2\frac{1+\zeta_A}{1+2\zeta_A} < 2. \tag{E.23}$$

$\square$