The Return of Consumption

Harjoat Bhamra* Peter Vaporakis† Paul Whelan‡

ABSTRACT

Under recursive preferences, the stochastic discount factor depends on the return on aggregate wealth, an unobservable encompassing human capital, housing, and private business. We derive closed-form expressions for the bias in risk aversion and the elasticity of intertemporal substitution that arises from proxying this return with a public equity index. The resulting distortions quantitatively reproduce the pathological estimates in extant empirical work. A joint estimator that recovers preferences and the tradable wealth share eliminates over 90 percent of the distortion. Decomposing the bias, we find that a single quantity, the non-tradable risk premium, accounts for most of the correction.

Keywords: Consumption, asset pricing, recursive preferences, the equity premium puzzle, the risk-free rate puzzle.

JEL classification: C13, C58, D81, E21, G12.

This Draft: March 2026

*Imperial College Business School. Email: h.bhamra@imperial.ac.uk
†Copenhagen Business School. Email: pv.fi@cbs.dk
‡Chinese University of Hong Kong. Email: paulwhelan@cuhk.edu.hk

The empirical failure of consumption-based asset pricing is finance's most enduring puzzle. Despite its theoretical elegance, the Consumption Capital Asset Pricing Model (CCAPM) produces implausibly high estimates of risk aversion and near-zero elasticity of intertemporal substitution when confronted with data. This paper demonstrates that a measurement problem – the unobservability of the return on aggregate wealth – is sufficient to explain these failures. We derive closed-form expressions for the resulting bias, validate them and a corrective joint estimator in a calibrated economy, and show how to recover preferences and the wealth share simultaneously.

Consumption growth is the fundamental determinant of the stochastic discount factor that prices all assets (Lucas, 1978; Breeden, 1979), yet confronting this theory with data has been a forty-year exercise in frustration. Hansen and Singleton (1983) first documented implausibly high risk aversion estimates. Mehra and Prescott (1985) crystallized the finding as the equity premium puzzle; Weil (1989) added the risk-free rate puzzle. The profession responded by enriching the model – recursive preferences (Epstein and Zin, 1989; Weil, 1990), habit formation (Campbell and Cochrane, 1999; Constantinides, 1990), long-run risks (Bansal and Yaron, 2004), disaster risk (Rietz, 1988; Barro, 2006) – or by refining consumption measurement: ultimate consumption (Parker and Julliard, 2005), fourth-quarter consumption (Jagannathan and Wang, 2007), unfiltered consumption (Kroencke, 2017), and garbage (Savov, 2011).

Figure 1 summarises the estimation issue: every point represents a nonlinear GMM estimate of risk aversion and the elasticity of intertemporal substitution from a different model–consumption measure pairs, and every point lands in the same pathological region: implausibly high risk aversion and near-zero elasticity of intertemporal substitution. No combination of preference specification or consumption refinement moves the estimates into the economically reasonable region. Under recursive utility, the stochastic discount factor depends jointly on consumption growth and the return on aggregate wealth. Aggregate wealth includes human capital, housing, and private business equity – none of which are directly observable. Empirical work, following Epstein and Zin (1991) and Yogo (2006), substitutes the stock market return for the unobservable wealth return. Roll (1977) showed that CAPM tests are joint tests of the model and the mean-variance efficiency of the market

proxy. We establish the analogous result for recursive preferences: estimates of risk aversion and the EIS inherit a systematic, quantifiable bias from the market proxy. Unlike Roll's critique, which is an impossibility result, ours admits a constructive resolution: we derive the bias in closed form and propose an estimator that corrects for it. The mechanism is specific to recursive preferences and vanishes under expected utility: the bias exists only when $\gamma \neq 1/\psi$.

We make four contributions. First, we provide an analytical characterization of the bias that arises from wealth mismeasurement. The bias operates through three distinct channels – the variance of the stochastic discount factor, the risk premium on the market return, and the level of the risk-free rate – each of which can be analyzed in isolation. The central insight is that these channels interact nonlinearly: their combined effect when the econometrician jointly fits all standard pricing restrictions is far larger than any single channel alone predicts.

The quantitative implications are large. When the econometrician substitutes the market return for the wealth return and estimates the model by standard GMM, the pseudo-true risk aversion rises by an order of magnitude and the EIS collapses toward zero. These are not finite-sample artifacts: they are population-level distortions – the values to which any consistent estimator converges. The distortion is nearly invariant to preferences: across a wide range of plausible calibrations, the pseudo-true risk aversion falls in a narrow band far from the truth. Estimates of risk aversion in the range commonly reported in the literature are therefore nearly uninformative about true preferences.

Second, we study the empirical properties of our analytical results via a Monte Carlo simulation. In a calibrated long-run risk economy (Bansal, Kiku, and Yaron, 2007), GMM estimation with the true return on wealth recovers the calibrated preference parameters. Replacing the wealth return with the market return inflates the median risk-aversion estimate more than sixfold and compresses the median EIS by a factor of forty. The theoretical predictions track the Monte Carlo results.

Third, we propose a joint estimator that recovers preferences and the tradable wealth share simultaneously within a standard GMM framework. Euler equation restrictions across multiple test assets jointly identify preference parameters and wealth composition, even when no single moment condition suffices. In simulation, this estimator eliminates over 90 percent

of the bias in long samples; with realistic sample lengths using *monthly* data our estimator displays good power and size properties, but *annual time-aggregated* estimation suffers severe finite-sample distortions.

Fourth, we connect our bias analysis to the "dark matter" framework of Chen, Dou, and Kogan (2024). CDK show that unknown preference parameters inflate estimation variance by a factor of $\varrho = 196$ in the BKY model. Decomposing the equilibrium Jacobian through the rank-two influence function to identify which dark matter components drive the *bias* from wealth mismeasurement, we find that the efficiency and bias rankings diverge sharply at the extremes. CDK's most informationally costly parameter – volatility persistence – contributes negligibly to the bias, while the dominant bias driver – long-run risk persistence – sits in the well-identified part of CDK's parameter space. At the level of non-tradable return moments, the non-tradable risk premium alone accounts for 91 percent of the bias reduction. The dark matter that matters most for preference estimation is the dark matter that is least dark.

Related Literature. Roll (1977) showed that CAPM tests are joint tests of the model and the mean-variance efficiency of the market proxy – an impossibility result, since the true market portfolio is unobservable. Shanken (1987) derived bounds on the resulting test distortions. We establish the analogous result for recursive-utility Euler equations, with a constructive difference: closed-form bias expressions and a joint estimator that recovers preferences and the wealth share. Several papers address the gap between market and total wealth returns: Lettau and Ludvigson (2001) use the consumption-wealth ratio as a conditioning variable, Lustig, Nieuwerburgh, and Verdelhan (2013) construct total wealth returns, Chen, Favilukis, and Ludvigson (2013) bypass the R_w proxy via semiparametric methods, and Guo and Savickas (2019) show that the choice of wealth proxy significantly affects the intertemporal risk-return tradeoff. We complement these by providing a quantitative assessment of the size and direction of biases induced in preference estimation by the wealth proxy problem.

Under Epstein–Zin preferences the SDF depends on R_w through $R_{w,t+1}^{\theta-1}$; under power utility it does not. The wealth return is therefore "dark matter" in the sense of Chen, Dou,

and Kogan (2024) – an unobservable component that degrades inference whenever $\psi \neq 1/\gamma$. CDK quantify the information cost of this dark matter, showing that unknown preferences inflate estimation variance by a factor of $\varrho = 196$ in the BKY model – a relative penalty on top of the absolute data requirements identified by Merton (1980). We decompose the equilibrium Jacobian through the rank-two influence function to identify which dark matter components drive the *bias* from wealth mismeasurement, and find that the bias-relevant dark matter lies in the well-identified part of CDK's parameter space. Ghosh, Julliard, and Taylor (2017) decompose the SDF into observable and unobservable components and show that the missing piece is large and countercyclical – consistent with the mechanism we formalize. Improved consumption data – ultimate consumption (Parker and Julliard, 2005), fourth-quarter growth (Jagannathan and Wang, 2007), unfiltered NIPA data (Kroencke, 2017), garbage (Savov, 2011) – can sharpen inference on preferences, but our point is orthogonal: even with perfect consumption measurement, the unobservability of R_w creates a distinct and quantitatively dominant bias. Stock and Wright (2000) and Kleibergen (2009) document weak identification in consumption Euler equations. Our analysis addresses a logically distinct problem: an omitted-variable bias that shifts the pseudo-true parameter values regardless of identification strength.

Our decomposition $R_w = \lambda R_m + (1 - \lambda) R_h$ draws on a literature that consistently finds tradable market wealth to be a minority of the total: $\lambda \approx 0.1$–0.3 (Jagannathan and Wang, 1996), human capital at two-thirds or more of total wealth (Campbell, 1996), labor income shares that vary countercyclically and predict returns (Santos and Veronesi, 2006), and labor income claims at roughly 70 percent of aggregate wealth (Favilukis and Lin, 2016). Our baseline calibration adopts $\lambda = 0.30$, near the centre of these estimates.

The remainder of the paper is organized as follows. Section I presents the model and formalizes the wealth measurement problem. Section II derives the exact bias characterizations for risk aversion and the elasticity of intertemporal substitution, including first-order correction procedures. Section III provides simulation evidence from a calibrated long-run risk economy, including validation of the joint estimator. Section IV connects our bias analysis to the dark matter framework of Chen, Dou, and Kogan (2024), comparing efficiency and bias rankings of unobservable model components. Section V concludes. Appendix G

provides empirical estimation details for the motivating figure.

I. Model and the Wealth Measurement Problem

This section derives the stochastic discount factor under recursive utility, formalizes the wealth measurement problem, and introduces a joint estimator that recovers the wealth share λ alongside preference parameters.

A. Preferences and the Stochastic Discount Factor

Consider a representative investor with Kreps and Porteus (1978)–Epstein and Zin (1989)–Weil (1990) recursive preferences defined by the constant elasticity of substitution aggregator:

$$U_t = \left[(1-\beta)(C_t)^{\frac{1-\gamma}{\theta}} + \beta(\mathbb{E}_t[U_{t+1}^{1-\gamma}])^{1/\theta}\right]^{\frac{\theta}{1-\gamma}}, \tag{1}$$

where β is a time discount factor, γ is the coefficient of relative risk aversion (RRA), ψ is the elasticity of intertemporal substitution (EIS), and

$$\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}. \tag{2}$$

When $\gamma > 1/\psi$ ($\gamma < 1/\psi$), the agent prefers earlier (later) resolution of uncertainty. The agent maximizes lifetime utility subject to the budget constraint

$$W_{t+1} = (W_t - C_t)\, R_{w,t+1}, \tag{3}$$

where W_t is total wealth, C_t is aggregate consumption, and $R_{w,t+1}$ is the gross return on the claim to aggregate consumption.

Epstein and Zin (1989) show that recursive preferences lead to a generalization of the expected utility first-order condition. Combining the first-order conditions for consumption and portfolio choice, and eliminating the value-function coefficient via the budget constraint, yields the consumption Euler equation (EZ 1991, eq. 13).[1] In our notation, this takes the

[1] A self-contained derivation tracing each algebraic step from Epstein and Zin (1991) equations (11)–(12) to equation (13), together with simulation and estimation codes, is available in the companion note "Deriving the Epstein–Zin

form:

$$1 = \mathbb{E}_t \left[\underbrace{\beta^\theta \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\theta}{\psi}} R_{w,t+1}^{\theta-1}}_{M_{t+1}} R_{j,t+1} \right] = \mathbb{E}_t \left[M_{t+1} R_{j,t+1} \right], \tag{4}$$

where M_{t+1} denotes the stochastic discount factor (SDF). In log form, the SDF is:

$$m_{t+1} = \theta \log \beta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1) r_{w,t+1}. \tag{5}$$

The wealth Euler equation. A restriction that plays a central role in our joint estimator is obtained by setting $j = w$ (the wealth portfolio) in (4) and using $M_{t+1} R_{w,t+1} = \beta^\theta (C_{t+1}/C_t)^{-\theta/\psi} R_{w,t+1}^\theta$. Then, subtracting one and dividing by θ (following EZ 1991, eq. 18):

$$\mathbb{E}_t \left[\frac{1}{\theta} \left(\left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-1/\psi} R_{w,t+1} \right)^\theta - 1 \right) \right] = 0. \tag{6}$$

This is an Euler equation for the wealth portfolio that does not involve any additional asset return R_j –it restricts the joint dynamics of consumption growth and the wealth return alone. When $\theta \to 0$ (logarithmic risk preferences), the expression converges via L'Hôpital's rule to $\log \beta - (1/\psi)\mathbb{E}_t[\Delta c_{t+1}] + \mathbb{E}_t[r_{w,t+1}] = 0$, which pins down ψ from the consumption-growth/wealth-return tradeoff even when risk aversion is logarithmic.

Key observation. The SDF depends on $r_{w,t+1}$, which is *unobservable*. A large component of aggregate wealth –human capital, housing, private business equity –is non-tradable, so neither market prices nor returns are directly observed. This is the root of the measurement problem analyzed below.

The canonical CCAPM with power utility is the special case $\gamma = 1/\psi$ (i.e., $\theta = 1$), yielding:

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}, \tag{7}$$

a one-factor model that does not depend on the return on wealth. The general class of recursive preferences implies a factor structure driven by shocks to both consumption *and* the return on wealth.

(1991) Euler Equation" (Bhamra, Vaporakis, and Whelan, 2026).

B. Wealth Return Decomposition

In applied work, the return on wealth is commonly approximated by the return on the aggregate stock market $R_{m,t}$, following Epstein and Zin (1991), Stock and Wright (2000), Weber (2000), and Yogo (2006), among others.

Assumption 1 (Wealth Return Structure)**:** *The log aggregate wealth return admits the portfolio decomposition:*

$$r_{w,t+1} = \lambda\, r_{m,t+1} + (1-\lambda)\, r_{z,t+1}, \tag{8}$$

where:

(i) $\lambda \in (0,1)$ *is the wealth share of tradable assets,*

(ii) $Var(r_{z,t+1}) = \sigma_z^2$,

(iii) $Corr(r_{m,t+1}, r_{z,t+1}) = \rho_{mz}$.

Assumption 2 (Conditional Normality)**:** *The returns* $r_{m,t+1}$ *and* $r_{z,t+1}$ *are conditionally jointly normal with finite variances* $\sigma_m^2 = Var(r_{m,t+1})$ *and* $\sigma_z^2 = Var(r_{z,t+1})$*, and consumption growth* Δc_{t+1} *has finite second moments.*

The parameter λ captures the share of publicly traded equity in total wealth. Estimates from the Flow of Funds suggest $\lambda \approx 0.1$–0.3 (Jagannathan and Wang, 1996; Lustig, Nieuwerburgh, and Verdelhan, 2013). The component $r_{z,t+1}$ represents returns to human capital, housing, private equity, and other non-tradable assets. The correlation ρ_{mz} between equity and non-tradable wealth returns is left unrestricted. When $\rho_{mz} = 0$, the components are orthogonal and the analytical results simplify. One expects $\rho_{mz} > 0$ since both equity and human capital returns are procyclical.

C. The Econometrician's Problem

Denote the preference parameters by $\boldsymbol{\vartheta} = (\gamma, \psi)$. The true moment condition for pricing asset j is:

$$g(\boldsymbol{\vartheta}; R_w) = \mathbb{E}\left[\beta^{\theta}\left(\frac{C_{t+1}}{C_t}\right)^{-\theta/\psi} R_{w,t+1}^{\theta-1} R_{j,t+1} - 1\right] = 0, \tag{9}$$

where $\theta = (1-\gamma)/(1-1/\psi)$ as above. Since $R_{w,t+1}$ is unobservable, the econometrician substitutes the market return and solves the pseudo-moment condition:

$$\tilde{g}(\boldsymbol{\vartheta}^*; R_m) = \mathbb{E}\left[\beta^{\theta^*}\left(\frac{C_{t+1}}{C_t}\right)^{-\theta^*/\psi^*} R_{m,t+1}^{\theta^*-1} R_{j,t+1} - 1\right] = 0. \tag{10}$$

Definition 1 (Pseudo-True Parameters)**.** The pseudo-true parameters $\boldsymbol{\vartheta}^* = (\gamma^*, \psi^*)$ solve:

$$\boldsymbol{\vartheta}^* = \arg\min_{\tilde{\boldsymbol{\vartheta}} \in \Theta} \|\tilde{g}(\tilde{\boldsymbol{\vartheta}}; R_m)\|_W^2 \tag{11}$$

for a positive definite weighting matrix W.

The pseudo-true parameter vector $\boldsymbol{\vartheta}^*$ is the probability limit of the GMM estimator when moment conditions are misspecified –the best approximation GMM can deliver given incorrect inputs. The concept originates in White (1982) for maximum likelihood estimation under misspecification; Hall and Inoue (2003) extend it to the GMM setting. The pseudo-true parameters are deterministic functions of the true parameters and λ; Section II characterizes this mapping exactly.

D. Estimation

The simulation study of Section III estimates preference parameters via Hansen and Singleton (1982)-style moment conditions:

$$\mathbb{E}\begin{bmatrix} M_t R_t^f - 1 \\ M_t R_{m,t}^e \\ M_t \mathbf{R_t^e} \end{bmatrix} = \mathbf{0}, \tag{12}$$

which jointly test the model's ability to price the risk-free rate, equity premium, and cross-section. We treat the discount factor β as known, calibrated to the relevant frequency ($\beta = 0.9989$ per month in the simulation study, following Bansal, Kiku, and Yaron 2007), and estimate using one-step GMM with identity weighting to avoid the well-known pitfalls of efficient GMM in this setting (Cochrane, 2005; Ludvigson, 2013).

E. Joint Estimation of the Wealth Share

The bias results of Section II depend on the wealth share λ, which so far enters as an external calibration parameter. This subsection shows that λ can be estimated jointly with preferences by exploiting observable proxies for non-tradable wealth returns.

E.1. Observable Wealth Return Decomposition

Assumption 3 (Portfolio Wealth Decomposition)**:** *The gross return on aggregate wealth admits the portfolio representation:*

$$R_{w,t+1} = \lambda\, R_{m,t+1} + (1-\lambda)\, R_{h,t+1}, \tag{13}$$

where $R_{m,t+1}$ is the gross return on publicly traded equity, $R_{h,t+1}$ is the gross return on non-tradable wealth, and $\lambda \in (0,1)$ is the share of tradable wealth in total wealth.

Here $R_{h,t+1}$ is the gross-return counterpart of $r_{z,t+1}$ from Assumption 1; we use h (for human capital) when working with gross returns. The decomposition follows from writing total wealth as $W_t = W_t^F + W_t^H$, where W_t^F is financial wealth and W_t^H is human capital, so that $\lambda = W_t^F/W_t$:

$$R_{w,t+1} = \frac{W_t^F}{W_t} R_{m,t+1} + \frac{W_t^H}{W_t} R_{h,t+1} = \lambda\, R_{m,t+1} + (1-\lambda)\, R_{h,t+1}. \tag{14}$$

We treat λ as constant, which holds exactly when wealth shares are stationary (as in the balanced-growth economy of Section III); more generally, λ represents the time-averaged wealth share $\mathbb{E}[W_t^F/W_t]$. The decomposition is standard in the conditional CAPM (Jagannathan and Wang, 1996) and consumption-based (Lettau and Ludvigson, 2001) literatures.

Remark 1 (Log vs. Gross Returns)**.** The λ in Assumption 3 is the wealth share W_t^F/W_t, which has a direct economic interpretation. The λ in Assumption 1 is the coefficient in the log-return decomposition $r_w = \lambda r_m + (1-\lambda) r_z$. These two λ's are not identical, but they coincide to first order.

First-order equivalence. Expanding $r_w = \log[\lambda e^{r_m} + (1-\lambda)e^{r_h}]$ around zero:

$$r_w = \lambda r_m + (1-\lambda) r_h + O(\sigma^2), \tag{15}$$

where the remainder is bounded by $\lambda(1-\lambda)(r_m - r_h)^2/2$. At monthly frequency ($\sigma_m \approx 5$ percent), the approximation error is below 0.1 percent –two orders of magnitude smaller than the bias effects documented in Section II.

Convergence of the Taylor expansion. To expand the gross return $R_w = \lambda R_m + (1-\lambda) R_h$ in a power series, one factors out the larger term. When $\lambda R_m > (1-\lambda) R_h$, one writes $R_w = \lambda R_m [1 + (1-\lambda) R_h / (\lambda R_m)]$ and expands in the ratio $(1-\lambda) R_h / (\lambda R_m)$, which converges when this ratio has absolute value less than one. Which term dominates depends on the calibration: $\mathbb{E}[R_m] > \mathbb{E}[R_h]$ (equities earn a higher expected return than human capital, which behaves like a bond) but $\lambda < 1 - \lambda$ (human capital is the larger share of total wealth). With $\lambda \approx 0.3$, $\mathbb{E}[R_m] \approx 1.07$, and $\mathbb{E}[R_h] \approx 1.03$, we have $\lambda \mathbb{E}[R_m] \approx 0.32$ versus $(1-\lambda)\mathbb{E}[R_h] \approx 0.72$, so the non-tradable component dominates in expectation –human capital behaves like a low-beta claim with modest expected return but constituting the bulk of total wealth.

E.2. The Augmented SDF

Substituting the portfolio decomposition (13) into the Euler equation (4) yields the augmented SDF:

$$M_{t+1}(\gamma, \psi, \lambda) = \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-1/\psi} \left(\lambda \, R_{m,t+1} + (1-\lambda) \, R_{h,t+1} \right)^{1-1/\theta} \right]^{\theta}, \tag{16}$$

The SDF now depends on three estimable parameters: risk aversion γ, the EIS ψ, and the tradable wealth share λ.

Two limiting cases clarify the role of λ:

- $\lambda = 1$: The augmented SDF reduces to the standard specification with $R_w = R_m$, recovering the pseudo-true parameters of Section C. The joint estimator thus nests the conventional approach as a testable restriction.

- $\lambda = 0$: All wealth is non-tradable and $R_w = R_h$. The market return drops out of the SDF entirely.

E.3. Augmented Moment Conditions

The augmented Euler equations are:

$$\mathbb{E}\begin{bmatrix} M_t(\gamma,\psi,\lambda)\, R_t^f - 1 \\ M_t(\gamma,\psi,\lambda)\, R_{m,t}^e \\ M_t(\gamma,\psi,\lambda)\, \mathbf{R_t^e} \\ \frac{1}{\theta}\left[\left(\beta\, g_t^{-1/\psi}\, \hat{R}_{w,t}\right)^{\theta} - 1\right] \end{bmatrix} = \mathbf{0}, \tag{17}$$

where $g_t = C_{t+1}/C_t$ denotes gross consumption growth and $\hat{R}_{w,t} = \lambda R_{m,t} + (1-\lambda)R_{h,t}$ is the constructed wealth return.

Origin and role of the fourth condition. The fourth moment condition is the wealth Euler equation (6), derived from Epstein and Zin (1991) equation 13 via their equation 18 (see Section A), with $R_{w,t}$ replaced by the constructed return $\hat{R}_{w,t} = \lambda R_{m,t} + (1-\lambda)R_{h,t}$. It restricts the joint dynamics of consumption growth and the wealth return *without* involving any additional test asset return $R_{j,t}$. This condition provides the primary source of identification for λ: it requires that the SDF correctly prices its own wealth return, pinning down the tradable wealth share that reconciles the consumption-growth and wealth-return dynamics. The division by θ ensures the condition is well-scaled as θ varies and delivers a well-defined limiting moment as $\theta \to 0$.

The fourth condition is linearly independent of the first three because it prices a different portfolio. Conditions 1–3 price R^f, $R_m - R^f$, and $\mathbf{R}_j - R^f$ respectively; the fourth prices the *constructed* wealth return $\hat{R}_w = \lambda R_m + (1-\lambda)R_h$, which is a distinct asset whenever $R_h \neq R_m$ (i.e., whenever $\lambda < 1$). The key distinction is that $\hat{R}_w$ depends on λ through both its composition and the SDF's θ-power, whereas the market and cross-sectional conditions do not.

With the 25 Fama and French (1992) size and book-to-market portfolios, the system provides $N+3 = 28$ moment conditions for 3 parameters, yielding 25 overidentifying restric-

tions.

E.4. Identification of λ

Proposition 1 (Identification of λ)**:** *The parameter vector (γ, ψ, λ) is locally identified from the moment conditions* (17) *provided:*

(i) $R_{m,t+1} \neq R_{h,t+1}$ *with positive probability (exclusion),*

(ii) $\gamma \neq 1/\psi$ *(i.e.,* $\theta \neq 1$*; recursive preferences) and* $\psi \neq 1$*, and*

(iii) the Jacobian $\nabla_{(\gamma,\psi,\lambda)}\, g(\gamma, \psi, \lambda)$ *has full column rank at the true parameter values (rank condition).*

The proof is in Appendix A.

Intuition. If λ is set too high, the constructed wealth return over-weights equities and the SDF misprices assets whose returns load heavily on human capital risk. If λ is set too low, the SDF misprices assets with high equity market exposure. The cross-section of test portfolios –with heterogeneous exposures to both market and labor income risk –pins down the locally identified λ that reconciles all Euler equations simultaneously.

This cross-sectional role is fundamentally different from the baseline estimator's use of the cross-section. The baseline substitutes R_m for R_w, misspecifying the SDF; adding cross-sectional moments compounds the misspecification by requiring the distorted SDF to price more assets. The joint estimator, by contrast, embeds λ in a correctly specified SDF, so cross-sectional moments provide additional identifying restrictions for λ without introducing misspecification. Section D confirms that expanding from 5 to 25 test assets tightens the joint estimates rather than distorting them.

II. Main Theoretical Results

This section derives the analytical consequences of substituting r_m for r_w in the Epstein-Zin Euler equation. We proceed in five steps: (i) a closed-form variance-matching characterization of the risk-aversion bias (Theorem 2), (ii) the EIS distortion (Theorem 3), (iii) the full Euler equation pseudo-true, which shows that multiple bias channels interact multiplicatively

(Proposition 4), (iv) sensitivity analysis (Section D), and (v) implementable bias-correction procedures (Section E).

A. Variance-Matching Bias Characterization

We begin by isolating the SDF *variance* channel: the pseudo-true γ^* that equates $\text{Var}(m^*)$ under the misspecified model to $\text{Var}(m)$ under the truth, holding ψ at the true value. This partial characterization –which abstracts from risk premium and risk-free rate matching (addressed in Proposition 4) –admits an exact, non-perturbative solution.

Theorem 2 (Variance-Matching Bias Characterization)**:** *Under Assumptions 1 and 2, the* γ^* *that equates* $Var(m^*(\gamma^*, \psi; r_m))$ *to* $Var(m(\gamma, \psi; r_w))$ *–holding the EIS at the true* ψ *– satisfies:*

$$\gamma^* = \frac{1}{\psi} + \left(\gamma - \frac{1}{\psi}\right)\sqrt{\lambda^2 + (1-\lambda)^2\frac{\sigma_z^2}{\sigma_m^2} + 2\lambda(1-\lambda)\rho_{mz}\frac{\sigma_z}{\sigma_m}} \tag{18}$$

where $\sigma_m^2 = Var(r_{m,t+1})$*,* $\sigma_z^2 = Var(r_{z,t+1})$*, and* $\rho_{mz} = Corr(r_{m,t+1}, r_{z,t+1})$*.*

The proof is in Appendix A.

Remark 2. Theorem 2 characterizes a *variance-matching* pseudo-true, which is generally distinct from the GMM pseudo-true of Definition 1. Its value is analytic transparency: it isolates the SDF variance channel and yields a closed-form solution, making the mechanism clear. As Proposition 4 shows below, additionally imposing risk premium and risk-free rate matching amplifies the bias dramatically.

Define the scaling factor $\rho \equiv \sqrt{\lambda^2 + (1-\lambda)^2\sigma_z^2/\sigma_m^2 + 2\lambda(1-\lambda)\rho_{mz}\sigma_z/\sigma_m}$. Then $\gamma^* = 1/\psi + \rho(\gamma - 1/\psi)$: the pseudo-true γ^* lies on the ray from $1/\psi$ toward γ, scaled by ρ. When $\rho = \lambda$ (i.e., $\sigma_z = 0$), the formula recovers the zeroth-order special case $\gamma_0^* = \lambda\gamma + (1-\lambda)/\psi$ –a convex combination of risk aversion and the inverse EIS.

Three cases arise, governed by the preference wedge $\gamma - 1/\psi$:

- $\gamma > 1/\psi$ (preference for early resolution –the empirically relevant case). The zeroth-order term ($\sigma_z = 0$) attenuates γ^* below γ, but unobserved wealth variance ($\sigma_z^2 > 0$) inflates ρ above λ, pushing γ^* back up and potentially well above γ.

- $\gamma = 1/\psi$ (time-additive expected utility, $\theta = 1$). The SDF does not depend on the wealth return, so there is *no bias* regardless of λ, σ_z^2/σ_m^2, or ρ_{mz}.

- $\gamma < 1/\psi$ (preference for late resolution). The bias works in the opposite direction.

The second case is sharp: the measurement problem is specific to recursive preferences and vanishes under expected utility.

B. Bias in the Elasticity of Intertemporal Substitution

Theorem 2 holds ψ fixed and characterizes only the risk-aversion bias. We now ask: what happens to the EIS when the full SDF is misspecified?

Theorem 3 (EIS Bias)**:** *Under Assumptions 1 and 2, the pseudo-true EIS satisfies, to first order in* $(1-\lambda)$*:*

$$\frac{1}{\psi^*} = \frac{\lambda}{\psi}\left(1 + \frac{(1-\lambda)\gamma\rho_{cm}\sigma_c}{\lambda\sigma_m}\right) + (1-\lambda)\gamma\rho_{cz}\frac{\sigma_z}{\sigma_c} + O((1-\lambda)^2), \tag{19}$$

where $\sigma_c = Std(\Delta c_{t+1})$*,* $\rho_{cm} = Corr(\Delta c_{t+1}, r_{m,t+1})$*, and* $\rho_{cz} = Corr(\Delta c_{t+1}, r_{z,t+1})$*.*

The variance channel of Theorem 2 alone implies $\psi^* = \psi$ –no EIS bias (proved in Appendix A). The EIS distortion arises entirely from the risk premium channel: lacking the unobserved component r_z, the consumption loading θ^*/ψ^* must absorb the missing risk premium, inflating $1/\psi^*$ and compressing the EIS toward zero.

C. Full Euler Equation Pseudo-True

Theorems 2 and 3 isolate individual channels. We now impose the complete set of Euler equation restrictions simultaneously.

Proposition 4 (Full Euler Pseudo-True)**:** *Under joint log-normality, the pseudo-true parameters* (γ^*, ψ^*) *from the misspecified Euler equation* $\mathbb{E}[\exp(m_{t+1}^* + r_{j,t+1})] = 1$ *for* $j \in \{f, m\}$*, where*

$$m^* = \theta^* \log\beta - \frac{\theta^*}{\psi^*}\Delta c_{t+1} + (\theta^* - 1)r_{m,t+1}, \qquad \theta^* = \frac{1-\gamma^*}{1-1/\psi^*}, \tag{20}$$

solve the 2×2 *system:*

$$\mathbb{E}[m^* + r_f] + \tfrac{1}{2} Var(m^* + r_f) = 0 \tag{21}$$

$$\mathbb{E}[m^* + r_m] + \tfrac{1}{2} Var(m^* + r_m) = 0, \tag{22}$$

where all expectations and variances are evaluated under the true data-generating process. Assume further that $\gamma \in [5, 20]$*,* $\psi \in [1, 3]$*, and* $\lambda \in (0.1, 0.9)$*. We verify numerically that the system has a unique solution satisfying* $\gamma^*_{FE} > \gamma$ *and* $\psi^*_{FE} < \psi$*.*

Decomposition into three channels. The bias decomposes into three channels that interact nonlinearly through $\theta^* = (1 - \gamma^*)/(1 - 1/\psi^*)$:

1. **SDF variance channel** (Theorem 2): Matching $\text{Var}(m^*)$ alone gives $\gamma^*_{\text{VM}} = 1/\psi + (\gamma - 1/\psi)\rho$ and $\psi^* = \psi$ –no EIS distortion.

2. **Risk premium channel** (Theorem 3): Since $\text{Cov}(m, r_m)$ depends on $\text{Cov}(r_w, r_m) \neq \text{Var}(r_m)$, the consumption loading θ^*/ψ^* must absorb the discrepancy, distorting ψ^* toward zero.

3. **Risk-free rate channel**: The level constraint forces a joint recalibration of both γ^* and ψ^*.

Imposing all three conditions simultaneously amplifies the bias far beyond what any single channel predicts. At the Bansal, Kiku, and Yaron (2007) calibration ($\sigma(r_w)/\sigma(r_m) = 0.16$), variance matching alone gives $\gamma^*_{\text{VM}} = 2.2$; the full Euler system gives $\gamma^*_{\text{FE}} = 102.7$ and $\psi^*_{\text{FE}} = 0.04$.

D. Sensitivity to True Parameters

A natural concern is whether these predictions are an artifact of the specific BKY calibration. Figure 4 plots γ^*_{FE} across a grid of true (γ, ψ) pairs: γ^*_{FE} varies only between 98 and 111 as the true γ ranges from 5 to 20, while ψ^*_{FE} universally collapses to near zero. This near-invariance implies that observed estimates of $\gamma \approx 70$–100 are essentially uninformative about the true risk aversion: the entire region of plausible true parameters maps to a narrow band of pathological pseudo-true estimates.

E. Bias Correction

The theoretical results above characterize the bias; we now show how to undo it.

E.1. Risk Aversion

Setting $\sigma_z = 0$ in Theorem 2 yields the zeroth-order special case $\gamma_0^* = \lambda\gamma + (1-\lambda)/\psi$ (Corollary 7 in the Appendix), which expresses the pseudo-true as a convex combination of γ and $1/\psi$. Rearranging:

$$\gamma_0^* - 1 = \lambda(\gamma - 1) + (1-\lambda)\left(\frac{1}{\psi} - 1\right). \tag{23}$$

When $\gamma \gg 1/\psi$, the second term is negligible and the dominant scaling is $\gamma_0^* \approx \lambda\gamma$, yielding a simple deflation-based correction.

Corollary 5 (First-Order Corrected Estimator)**:** *Given a consistent estimator $\hat{\lambda}$ of the wealth share, the first-order bias-corrected risk aversion estimator is:*

$$\hat{\gamma}_{BC} = \hat{\lambda}(\hat{\gamma}_{GMM} - 1) + 1. \tag{24}$$

By the delta method, $\mathrm{Var}(\hat{\gamma}_{BC}) \approx \hat{\lambda}^2\mathrm{Var}(\hat{\gamma}_{GMM}) + (\hat{\gamma}_{GMM} - 1)^2\mathrm{Var}(\hat{\lambda}) + 2\hat{\lambda}(\hat{\gamma}_{GMM} - 1)\mathrm{Cov}(\hat{\gamma}_{GMM}, \hat{\lambda})$, where the covariance term vanishes when $\hat{\lambda}$ is estimated from an independent data source.

E.2. Elasticity of Intertemporal Substitution

Setting the cross-correlations ρ_{cm} and ρ_{cz} to zero in Theorem 3 gives $1/\psi^* \approx \lambda/\psi$, so:

$$\hat{\psi}_{BC} = \hat{\psi}^*/\hat{\lambda}. \tag{25}$$

The zero-correlation approximation is strongest when consumption growth has low correlation with financial returns, as is the case in postwar U.S. data.[2]

[2] In the BKY calibration, the joint estimator of Section E –which does not rely on this approximation –recovers ψ more reliably than the deflation formula; see Section D.

E.3. Sensitivity to the Wealth Share λ

Both corrections require a value for λ. Rather than commit to a single estimate, we present the bias correction as a function of λ across the range implied by the literature.

Estimates of the tradable wealth share span a wide interval. From Flow of Funds data, Jagannathan and Wang (1996) report $\lambda \approx 0.10$–0.30; Campbell (1996) estimates financial wealth at roughly one-third of total wealth ($\lambda \approx 0.33$); Santos and Veronesi (2006) calibrate the complement $(1 - \lambda)$ to the aggregate labor share, implying $\lambda \approx 0.30$; and Favilukis and Lin (2016) estimate $\lambda \approx 0.30$ in a structural model. The Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019) historical returns on all major asset classes (1870–2015) imply $\lambda \approx 0.48$.

Table VI reports the bias-corrected risk aversion $\hat{\gamma}_{BC}$ and EIS $\hat{\psi}_{BC}$ for λ ranging from 0.10 to 0.70. The corrected $\hat{\gamma}_{BC}$ is monotonically increasing in λ: at the lower end of the literature range ($\lambda = 0.15$), the median correction yields $\hat{\gamma}_{BC} \approx 12$; at the upper end ($\lambda = 0.50$), it yields $\hat{\gamma}_{BC} \approx 37$. The identified set narrows substantially as λ is pinned down more precisely –the interval $\lambda \in [0.20, 0.40]$ maps to $\hat{\gamma}_{BC} \in [16, 30]$, a range that excludes the pathological estimates ($\gamma^* > 70$) obtained under the market proxy.[3]

For a data-driven approach that does not require an external λ estimate, Section E introduces a joint estimator that recovers λ alongside preference parameters from the augmented moment conditions.

Section III evaluates these theoretical predictions –including the bias magnitudes, the three-channel decomposition, and the correction formulas –against Monte Carlo evidence from a calibrated long-run risk economy.

III. Simulation Study

The theoretical results of Section II make sharp quantitative predictions: replacing the return on wealth with the market return should inflate risk aversion severalfold and compress the EIS toward zero. We now test these predictions in a calibrated long-run risk economy with

[3]In the BKY calibration used for simulation (Section III), the model-implied λ depends on the volatility ratio $\sigma(r_w)/\sigma(r_m)$. The simulation exercises evaluate robustness across $\lambda_0 \in \{0.12, 0.30, 0.50\}$.

stochastic volatility, multiple assets, and finite samples.

A. Design

The data-generating process is the long-run risks model of Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2007), with consumption dynamics

$$g_{c,t+1} = \mu_c + x_t + \sigma_t \eta_{t+1}, \tag{26}$$

$$x_{t+1} = \rho_x x_t + \varphi_e \sigma_t e_{t+1}, \tag{27}$$

$$\sigma^2_{t+1} = \bar{\sigma}^2 + \nu(\sigma^2_t - \bar{\sigma}^2) + \sigma_w w_{t+1}, \tag{28}$$

where $\eta_{t+1}, e_{t+1}, w_{t+1}$ are i.i.d. $N(0,1)$ and mutually orthogonal. The return on aggregate consumption wealth is $R_{w,t+1} = (P_{c,t+1} + C_{t+1})/P_{c,t}$. This object is computable from the model but unobservable in practice –which is precisely the measurement problem the theory addresses. Table I reports parameter values.[4] The true preference parameters are $\gamma = 10$ and $\psi = 1.5$.

We simulate 1,000 independent paths of 2,000 years of monthly data. For each path, we estimate preference parameters via one-step GMM (identity weighting, Newey-West standard errors) using two proxies for the wealth return: the true R_w and the market R_m.[5] The experiment isolates a single source of variation: the proxy for wealth. Preferences, consumption data, and econometric method are held constant across specifications. In the BKY economy, the wealth return is far less volatile than the market return ($\sigma(r_w)/\sigma(r_m) = 0.16$), so substituting R_m for R_w injects substantial spurious variation into the SDF.

Construction of the non-tradable return. The joint estimator requires the non-tradable return R_h as an input. In the simulation, the true return on consumption wealth

[4] A companion teaching note (Bhamra, Vaporakis, and Whelan, 2026, "Teaching Notes: Recursive Preferences and the Long-Run Risk Model") provides a self-contained derivation of the BKY model solution –including the consumption claim, SDF decomposition, risk prices, real bond prices, and equity valuation –together with numerical calibration examples and the simulation and estimation code. Available from the authors on request.

[5] The estimation uses three moment conditions –a risk-free rate Euler equation, a market excess return Euler equation, and the wealth Euler equation of (4) –giving one overidentifying restriction. The discount factor β is fixed at its calibrated value. Appendix F provides details on starting values, optimizer bounds, and convergence diagnostics.

R_w is known from the model solution, so we back out R_h from the portfolio identity:

$$R_{h,t+1} = \frac{R_{w,t+1} - \lambda_0 R_{m,t+1}}{1 - \lambda_0}, \tag{29}$$

where λ_0 is the structural wealth share. This construction makes the gross-return decomposition $R_w = \lambda R_m + (1 - \lambda) R_h$ hold exactly by definition. The simulation therefore validates the joint estimator under ideal conditions: the econometrician observes an R_h that is perfectly consistent with the true wealth decomposition. The log-return counterpart $r_w \approx \lambda r_m + (1 - \lambda) r_h$ introduces a Jensen's inequality error of order $O(\sigma^2)$; at monthly frequency in the BKY economy, this approximation error is below 0.1 percent (see Appendix G).

Test assets. The joint estimator requires cross-sectional assets with heterogeneous risk exposures. The BKY model generates a cross-section of dividend claims whose log dividend growth follows:

$$g_{d,j,t+1} = \mu_{d,j} + \varphi_j \, x_t + \pi_j \, \sigma_t \, \eta_{t+1} + \psi_{d,j} \, \sigma_t \, u_{j,t+1}, \tag{30}$$

where η_{t+1} is the consumption shock (shared across all assets) and $u_{j,t+1}$ is an idiosyncratic dividend shock. The five baseline claims differ along four dimensions: leverage on the long-run risk state ($\varphi_j \in [2.0, 5.4]$), sensitivity to the consumption shock ($\pi_j \in [1.5, 3.8]$), idiosyncratic volatility ($\psi_{d,j} \in [5.2, 7.3]$), and mean dividend drift ($\mu_{d,j}$). This heterogeneity generates cross-sectional variation in the covariance of returns with both R_m and R_h, which the augmented moment conditions exploit to identify λ. Appendix F shows that expanding to 25 dividend claims tightens all estimates.

B. Results

Table II reports the main findings. With the true return on wealth, the estimator is approximately unbiased: $\hat{\gamma}_{\text{med}} = 11.3$ and $\hat{\psi}_{\text{med}} = 1.65$, with cross-path standard deviations of 0.75 for $\hat{\gamma}$ and 0.07 for $\hat{\psi}$. The modest upward bias in γ arises because the BKY model is solved via log-linearization, whereas GMM evaluates exact level Euler equations; the approximation

error in model-implied moments translates into a small parameter shift.[6]

Replacing R_w with R_m dramatically distorts the estimates: $\hat{\gamma}_{\text{med}}$ rises from 11.3 to 73.7 (more than a sixfold increase) and $\hat{\psi}_{\text{med}}$ falls from 1.65 to 0.04 (a forty-fold compression). These distortions reflect a population-level bias, not sampling noise: the cross-path standard deviation of $\hat{\gamma}$ under R_m is 18.6, placing the true value $\gamma = 10$ more than three Monte Carlo standard deviations below the cross-path mean of 69.3.

Figures 2 and 3 decompose the bias by varying one design element at a time –wealth proxy, estimation frequency, and sample length:

1. *Wealth mismeasurement*: switching from R_w to R_m shifts $\hat{\gamma}_{\text{med}}$ from 11.3 to 73.7, accounting for the dominant share of the total distortion.

2. *Time aggregation*: estimating at annual rather than monthly frequency introduces substantial additional bias even with the true R_w, as $\hat{\gamma}_{\text{med}}$ rises from 11.3 (monthly) to 53.6 (annual).

3. *Small samples*: reducing from 2,000 to 62 years substantially inflates dispersion; in the correctly specified model (R_w), the median bias remains modest ($\hat{\gamma}_{\text{med}} = 14.1$ versus true $\gamma = 10$), far smaller than the wealth-mismeasurement effect.

Combining all three sources –annual frequency, market proxy, short sample –reproduces the central puzzles of the consumption-based asset pricing literature: estimated risk aversion far exceeding plausible priors (Mehra and Prescott, 1985) and near-zero EIS estimates (Hall, 1988).

C. Theory versus Simulation

Table III compares the analytical approximations of Section II with the simulation outcomes. The theoretical pseudo-true values γ^* are population quantities; the Monte Carlo estimates $\hat{\gamma}_{\text{med}}$ are their finite-sample counterparts. The hierarchy of predictions tracks the Monte Carlo results closely. Variance matching (Theorem 2) captures the direction of bias ($\gamma^* = 2.2$)

[6]Validation under exact Euler equations (discrete-state Markov chain) confirms that this residual bias is entirely attributable to the log-linear approximation; see Appendix F.

but understates its magnitude because it isolates only the SDF variance channel. The log-normal full Euler prediction (Proposition 4) overshoots at $\gamma^* = 102.7$, reflecting the Gaussian approximation. Numerically solving the level Euler equations from simulated BKY data gives $\gamma^* = 81.7$ –substantially closer to the Monte Carlo median of 73.7 than the log-normal prediction, with the remaining gap attributable to the non-Gaussian features of stochastic volatility. All three approximations predict $\psi^* < 0.05$, confirming that the EIS compression is robust across methods.

D. Joint Estimation of (γ, ψ, λ)

The preceding subsections establish that the measurement bias is large and theoretically well-understood. The joint estimator of Section E offers a constructive remedy. Table IV validates it within the BKY framework across three structural wealth shares ($\lambda_0 \in \{0.12, 0.30, 0.50\}$), estimating (γ, ψ, λ) jointly from the augmented moment conditions of Proposition 1 using five cross-sectional dividend claims from the BKY model.[7]

Risk aversion drops from $\hat{\gamma}_{\text{med}} = 71.6$ (market proxy, Table IV) to $\hat{\gamma}_{\text{med}} = 7.5$ (joint estimator, $\lambda_0 = 0.30$), eliminating over 90 percent of the bias. The EIS recovers from $\hat{\psi}_{\text{med}} = 0.04$ to 1.44, close to the true 1.5. The wealth share $\hat{\lambda}_{\text{med}}$ tracks the true λ_0 at all three calibrations (0.15 versus 0.12, 0.32 versus 0.30, 0.51 versus 0.50). The joint estimator undershoots the true $\gamma = 10$ by approximately 2.5 units –a residual attributable to the BKY log-linear approximation, as confirmed by exact-model validation in Appendix F. This remaining bias is an order of magnitude smaller than the market-proxy distortion.

Short-sample power. Table V addresses the practical question: does the estimator work with realistic sample sizes? With 62 years of *monthly* data ($T = 744$), the joint estimator yields $\hat{\gamma}_{\text{med}} = 6.4$ (versus 45.3 under the market proxy), $\hat{\psi}_{\text{med}} = 1.35$ (versus 0.05), and $\hat{\lambda}_{\text{med}} = 0.33$ (close to the true 0.30). A Wald test of $H_0 : \lambda = 1$ achieves 83 percent power at the 5 percent level, providing strong statistical evidence against equating market and wealth returns.

At *annual* frequency, by contrast, 62 observations are insufficient: the joint estimator

[7] Table XIII in Appendix F shows that expanding to 25 dividend claims tightens all estimates, confirming that the results are robust to cross-section size.

yields $\hat{\gamma}_{\text{med}} = 0$ (a boundary solution) and $\hat{\lambda}_{\text{med}} = 0.65$ (severely biased), and even the oracle (R_w) exhibits extreme noise (std = 45.9). Monthly data are therefore essential for empirical implementation.[8]

IV. Which Dark Matter Matters?

Chen, Dou, and Kogan (2024) develop a framework for quantifying the information cost of unobservable model components—"dark matter"—in structural estimation. Working within the Bansal, Kiku, and Yaron (2007) long-run risk model, they compute population Jacobians of observable moments with respect to all model parameters, partition these into structural and preference blocks, and derive information matrices that measure how much harder estimation becomes when preferences are treated as unknown nuisance parameters. Their headline result is $\varrho = 196$: the presence of unknown preference parameters inflates the variance of structural parameter estimates by a factor equivalent to needing nearly two centuries of additional data. The result is purely analytical—a property of the model's moment structure at the population level—and provides a powerful diagnostic for the fragility of inference in calibrated economies.[9]

We study the complementary question. CDK ask how preference uncertainty degrades inference on the economy; we ask how mismeasurement of a single economic variable—the return on aggregate wealth—distorts inference on preferences. CDK's pathology is one of *variance*: confidence intervals widen, but centre on the truth. Ours is one of *bias*: point estimates converge, but to the wrong values. Both analyses share the same mathematical building block—the GMM influence function $\Lambda = -(G'WG)^{-1}G'W$, introduced by Andrews, Gentzkow, and Shapiro (2017) as a sensitivity measure and implicit in CDK's selection matrix A—but apply it to different questions. CDK evaluate the full norm $\|G_{\eta,k}\|$ of each structural parameter's equilibrium effect across all moment directions; we pass the same object through Λ, a rank-two filter that selects only the directions relevant for preference estimation.

[8]The results are robust to alternative preference calibrations and to the number of cross-sectional test assets; see Appendix F for details.

[9]CDK work with the BKY 2012 calibration (their specification M1), which sets $\rho_x = 0.975$ and $\nu = 0.999$. Our paper uses the BKY 2007 calibration with $\rho_x = 0.982$ and $\nu = 0.98$. Seven of the thirteen structural parameters differ between the two calibrations, with the near-unit-root ν in CDK's specification likely inflating their $\varrho_\nu = 97.8$ relative to ours. All numerical comparisons with CDK's rankings should be read with this caveat.

The connection to Merton (1980) provides perspective on three distinct obstacles to learning from financial data. Merton shows that expected returns require several centuries of data to estimate precisely—an absolute limit on learning from time series. CDK show that nuisance parameters impose a multiplicative penalty: structural parameter estimates require 196× more data when preferences are unknown. We show that more data cannot help at all when the wrong return is used: with infinite data, the econometrician converges to $\gamma^* \approx 103$ rather than $\gamma = 10$. A practitioner facing all three problems simultaneously has wide confidence intervals centred on the wrong number.

A. Bias Sensitivity and the Chain Rule

Both CDK and this paper use Jacobians, but the derivatives measure fundamentally different objects. CDK differentiate the moment conditions $g(\theta)$ with respect to the full parameter vector θ *at the true values* θ_0:

$$D = \left.\frac{\partial g}{\partial \theta}\right|_{\theta_0}.$$

This Jacobian enters the Fisher information matrix $I = D'\Omega^{-1}D$, which governs the *variance* of parameter estimates. The perturbation is statistical: sample moments fluctuate around their true population values, and D determines how much those fluctuations shift $\hat{\theta}$.

Our object is different. We study how the *pseudo-true* parameters $\vartheta^* = (\gamma^*, \psi^*)$ respond to structural perturbations—changes in the economy's dark matter η that shift the population moments themselves. The pseudo-true ϑ^* is deterministic; there is no sampling uncertainty. The relevant Jacobian is the *bias sensitivity matrix*:

$$S \;=\; \frac{\partial \vartheta^*}{\partial \eta} \;\in\; \mathbb{R}^{2\times d_\eta}. \tag{31}$$

This decomposes via the chain rule (Proposition 11; proof in Appendix A):

$$S \;=\; \underbrace{\frac{\partial \vartheta^*}{\partial \text{mom}}}_{\Lambda\ \text{(model-free)}} \cdot \underbrace{\frac{\partial \text{mom}}{\partial \eta}}_{G_\eta\ \text{(model-specific)}}, \tag{32}$$

where $\Lambda \in \mathbb{R}^{2\times q}$ depends only on the quadratic structure of the log-normal Euler equations—

identical across all Epstein–Zin models—and $G_\eta \in \mathbb{R}^{q \times d_\eta}$ depends on how the specific equilibrium (BKY, habit, disasters) maps structural parameters to observable moments. The same influence function Λ governs both analyses, but CDK use it to measure which moment directions carry *information about preferences* (determining confidence interval width), while we use it to measure which moment perturbations shift the *wrong answer* (determining bias magnitude).

To aggregate across preference parameters, define the *bias dark matter measure*:

$$\rho_B(\eta_k) \;=\; S_k' \, V_\vartheta^{-1} \, S_k \, \mathrm{Var}(\eta_k), \tag{33}$$

where $S_k = \partial \vartheta^* / \partial \eta_k$ is the k-th column of S and V_ϑ is the asymptotic variance of $\hat{\vartheta}$ under misspecification. When $\rho_B(\eta_k) \gg 1$, uncertainty in η_k dominates sampling uncertainty—reducing dark matter matters more than collecting data. When $\rho_B(\eta_k) \ll 1$, more data helps more than better dark matter measurement. The aggregate measure $\rho_B = \lambda_{\max}(S \, \Sigma_\eta \, S' \, V_\vartheta^{-1})$ parallels CDK's $\varrho = \lambda_{\max}(I_F^{1/2} \, I_B^{-1} \, I_F^{1/2} - I)$: both are scale-free Rayleigh quotients, but ϱ measures the worst-case efficiency loss from dark matter while ρ_B measures the worst-case bias exposure.

We measure alignment between the two rankings with the Spearman rank correlation $\tau_{BE} = \mathrm{Spearman}(\{\rho_B(\eta_k)\}, \{\varrho_k\})$. If $\tau_{BE} \approx 1$, the hardest-to-identify parameters are also the most consequential for bias; if $\tau_{BE} \approx 0$, the rankings are orthogonal.

B. Complementary Rankings

Table VII compares CDK's efficiency ranking with our bias ranking. The two agree moderately overall ($\tau_{BE} = 0.64$) but diverge sharply at the extremes:

- CDK identify volatility persistence ν ($\varrho_\nu = 97.8$) and unconditional consumption volatility $\bar{\sigma}$ ($\varrho_{\bar{\sigma}} = 48.9$) as the most informationally costly dark matter—the parameters hardest to pin down from observable data.

- The parameter most consequential for bias is long-run risk persistence ρ ($\rho_B = 656$), which CDK rank fourth ($\varrho_\rho = 1.1$). Conversely, ν ranks fourth for bias ($\rho_B = 0.6$).

The relationship between the two rankings is governed by the *bias distortion factor*:

$$\mathcal{D}_k \;\equiv\; \frac{\rho_B(\eta_k)}{\varrho_k}, \tag{34}$$

which measures the fraction of parameter k's information cost that converts into estimation bias.[10] CDK's most informationally costly parameters (ν, $\bar{\sigma}$) enter the SDF through higher-order variance terms that appear symmetrically in both the true and misspecified Euler equations; their effects largely cancel when projected through Λ. The parameters that drive bias (ρ, μ_c) operate through first-order mean and persistence terms that enter asymmetrically when the market return substitutes for wealth. Table X confirms this pattern across the (γ, ψ) grid: ρ dominates the bias ranking at every calibration point where the BKY equilibrium exists.

C. Geometric Interpretation

The bias distortion factor $\mathcal{D}_k$ is governed by the matrix

$$W_B \;=\; \Lambda' \, V_\vartheta^{-1} \, \Lambda \;\in\; \mathbb{R}^{q \times q}, \tag{35}$$

which replaces CDK's Fisher information norm with a bias-weighted norm. This matrix has a key structural property: *its rank is exactly two* (Proposition 12; proof in Appendix B). Since $\Lambda \in \mathbb{R}^{2\times q}$ maps q observable moments to two preference parameters, W_B has a two-dimensional range—the *bias-relevant subspace*—and a $(q-2)$-dimensional null space—the *bias-neutral zone*. Any dark matter parameter whose equilibrium effect $G_{\eta,k}$ lies in the bias-neutral zone has $\mathcal{D}_k = 0$: informationally costly in CDK's sense, but irrelevant for bias. The rank-two constraint is model-free; it reflects the fact that there are only two preference parameters to distort.

Figure 5 illustrates the geometry. The singular value decomposition $\Lambda = U\Sigma V'$ completes the description (Corollary 13; proof in Appendix C). The columns of $V = [v_1, v_2] \in \mathbb{R}^{q\times 2}$ are the two *principal bias directions* in moment space; the singular values $\sigma_1 \geq \sigma_2 > 0$ are

[10] The factorisation $\rho_B = \varrho_k \times \mathcal{D}_k$ is diagnostic rather than structural: ϱ_k and ρ_B are defined in different moment systems (CDK's autocovariance moments vs. our Euler moments) and with different variance normalisations. The ratio $\mathcal{D}_k$ is a useful comparison tool but should not be interpreted as a multiplicative decomposition.

the amplification factors along each direction. Every equilibrium effect decomposes into a bias-relevant component and a bias-neutral component:

$$G_{\eta,k} = \underbrace{VV'\,G_{\eta,k}}_{w_k \text{ (survives)}} + \underbrace{(I - VV')\,G_{\eta,k}}_{w_k^{\perp} \text{ (annihilated)}}. \tag{36}$$

The influence function discards the bias-neutral component entirely: $\Lambda\, G_{\eta,k} = U\Sigma\, w_k$ where $w_k = V'G_{\eta,k} \in \mathbb{R}^2$. The squared bias sensitivity is therefore

$$\|\Lambda\, G_{\eta,k}\|^2 = \sigma_1^2\, w_{k,1}^2 + \sigma_2^2\, w_{k,2}^2. \tag{37}$$

For volatility persistence ν: $\|w_\nu^{\perp}\| \gg \|w_\nu\|$—the equilibrium effect is nearly orthogonal to the bias-relevant subspace, so CDK's large efficiency cost is almost entirely annihilated by the projection ($\mathcal{D}_\nu \approx 0.006$). For long-run risk persistence ρ: $\|w_\rho\| \gg \|w_\rho^{\perp}\|$—the equilibrium effect lies within the bias-relevant subspace, converting nearly all of its information cost into bias ($\mathcal{D}_\rho \approx 596$).

D. The Non-Tradable Risk Premium

Table VIII decomposes the dark matter into four moments of the non-tradable return r_z: the risk premium $\pi_z = \mathbb{E}[r_z] - r_f$, the variance σ_z^2, the consumption covariance σ_{cz}, and the risk-free rate covariance σ_{zf}.

The non-tradable risk premium π_z alone accounts for 91 percent of the bias reduction when the correction is applied to the no-correction baseline (Table IX). Starting from the no-correction baseline ($\gamma = 15.6$, $\psi = 2.4$), adding the π_z correction reduces the bias norm by 91 percent to $(\gamma, \psi) = (10.5, 1.47)$, close to the true values. The variance term σ_z^2, when added alone, actually *increases* the bias (-61% "reduction"), while σ_{zf} is negligible (0.1%).

This dominance of the mean channel reflects the asymmetric structure of the forward map. The γ bias is primarily a *variance* effect—it arises from the mismatch between $\mathrm{Var}(r_w)$ and $\mathrm{Var}(r_m)$ in the SDF—while the ψ bias is primarily a *mean* effect—it arises from the missing risk premium in the equity-premium Euler equation. The non-tradable risk premium π_z is the only single correction that simultaneously improves both γ and ψ: it shifts the

equity-premium equation's intercept (correcting ψ) without distorting the variance balance (preserving the γ correction from the quadratic coefficient).

E. Implications

CDK's framework quantifies the total information cost of dark matter across all moment directions. Our analysis decomposes that cost into the fraction that distorts preference estimates and the fraction that does not.

The decomposition identifies a constructive resolution. The dark matter that matters most for bias—the non-tradable risk premium π_z, which alone accounts for 91 percent of the bias reduction—is driven by parameters (ρ, μ_c) that sit in the well-identified part of CDK's parameter space ($\varrho_\rho = 1.1$, $\varrho_{\mu_c} = 1.0$). The dark matter that CDK identify as most informationally costly (ν, $\bar{\sigma}$, with ϱ values of 97.8 and 48.9) enters through variance channels that are nearly orthogonal to the bias-relevant subspace. The dark matter that matters most for preference estimation is, in CDK's metric, the dark matter that is least dark.

The practical implication is that the single most valuable piece of dark matter for correcting preference estimates is the expected excess return on non-tradable wealth. Estimates of this quantity are available at the aggregate level from measures of labour income growth and housing returns, and at a more granular level from household-level panel data. CDK's analysis prescribes more data; ours prescribes the right return.

V. Conclusion

Using market returns as a proxy for the unobservable return on aggregate wealth is sufficient to explain the pathological preference estimates –high risk aversion, near-zero EIS –that have defined the empirical record of consumption-based asset pricing for four decades. We derive the bias in closed form, show that it arises from the nonlinear interaction of three pricing channels under recursive preferences, and demonstrate that it is nearly invariant to the true preference parameters: across a wide range of plausible calibrations, the pseudo-true risk aversion clusters in a narrow band far from the truth. The mechanism is specific to recursive preferences and vanishes under expected utility.

Monte Carlo simulation in a calibrated long-run risk economy confirms that the analytical predictions accurately describe what standard estimators produce. A joint estimator that recovers preferences and the tradable wealth share simultaneously eliminates over 90 percent of the bias, with good power and size properties at monthly frequency.

A comparison with the dark matter framework of Chen, Dou, and Kogan (2024) reveals that the model parameters hardest to identify from observable data are not the parameters that drive the bias. The dominant bias driver –long-run risk persistence –sits in the well-identified part of CDK's parameter space, while CDK's most fragile parameters contribute negligibly to the pseudo-true distortion. The single most valuable piece of unobservable information for correcting preference estimates is the non-tradable risk premium, which alone accounts for 91 percent of the bias reduction. This is an optimistic finding: the part of dark matter that matters most for preference estimation is the part that is least dark.

Welfare calculations, policy prescriptions, and model calibrations that treat published risk aversion estimates as structural may rest on parameters that primarily reflect measurement error rather than genuine preferences. Empirical progress requires either improving the measurement of broad wealth returns or applying the bias correction and joint estimation procedures developed here. The consumption-based asset pricing framework is theoretically sound; its apparent empirical failures are a data-availability problem, not a model failure.

Parameter	Symbol	Value
Consumption Process		
Mean growth	μ_c	0.00150
Persistence	ρ_x	0.982
Volatility shock	φ_e	0.042
Unconditional volatility	$\bar{\sigma}$	0.0054
Vol-of-vol persistence	ν	0.98
Vol-of-vol shock	σ_w	0.0000068
Preference Parameters (True)		
Risk aversion	γ	10
Elasticity of substitution	ψ	1.5
Discount factor	β	0.9989

Table I: **Simulation Calibration: Bansal-Kiku-Yaron Model**
Monthly calibration parameters for the long-run risk model. Consumption process parameters from Bansal, Kiku, and Yaron (2007).

	γ			ψ		
	Mean	Median	Std	Mean	Median	Std
Long sample						
Return on wealth	11.30	11.29	0.75	1.66	1.65	0.07
Return on market	69.30	73.65	18.58	0.08	0.04	0.14
Short sample						
Return on wealth	25.88	14.12	22.11	1.12	1.46	0.67
Return on market	34.30	16.77	36.53	0.34	0.06	0.33

(a) Monthly frequency

	γ			ψ		
	Mean	Median	Std	Mean	Median	Std
Long, return on wealth	44.31	53.60	23.83	0.22	0.02	0.39
Short, return on market	48.46	47.76	48.20	0.23	0.03	0.30

(b) Annual frequency

Table II: **Simulation Results: Nonlinear GMM**
Monte Carlo simulation with 1,000 paths of 2,000 years each. True parameters: $\gamma = 10$, $\psi = 1.5$. Bansal-Kiku-Yaron long-run risk model. R_w: true return on wealth. R_m: market return proxy.

	True Value	Variance Matching	Log-Normal Euler	Level Euler (simulation)	Monte Carlo Median	Monte Carlo Mean (SE)
γ^*	10	2.2	102.7	81.7	73.7	69.3 (18.6)
ψ^*	1.5	1.500	0.041	0.044	0.04	0.08 (0.14)

Table III: **Theory vs. Simulation: Pseudo-True Parameters**
"Variance Matching": Theorem 2 evaluated at BKY $\sigma_{rw}/\sigma_m = 0.16$. "Log-Normal Euler": Proposition 4, solving the 2×2 system from analytical BKY moments. "Level Euler": numerically solving $\mathbb{E}[M^* R_j] = 1$ from 10,000 years of simulated BKY data. "Monte Carlo": median and mean across 1,000 paths $\times$ 2,000 years, standard errors in parentheses.

	$\hat{\gamma}$		$\hat{\psi}$		$\hat{\lambda}$	
Estimator	Mean	Median	Mean	Median	Mean	Median
True values	10.0	–	1.50	–	–	–
Standard + R_m (biased)	55.6	71.6	0.19	0.04	–	–
Standard + R_w (oracle)	11.3	11.3	1.66	1.65	–	–
Joint ($\lambda_0 = 0.12$)	7.0	7.4	1.39	1.43	0.17	0.15
Joint ($\lambda_0 = 0.30$)	7.3	7.5	1.43	1.44	0.33	0.32
Joint ($\lambda_0 = 0.50$)	7.6	7.6	1.45	1.45	0.51	0.51

Table IV: **Structural λ Identification**
Joint estimation of (γ, ψ, λ) across three structural wealth shares. 1,000 paths × 2,000 years, monthly data, one-step GMM with identity weighting and 5 BKY dividend claims. True parameters: $\gamma = 10$, $\psi = 1.5$.

	$\hat{\gamma}$			$\hat{\psi}$			$\hat{\lambda}$		
Estimator	Mean	Median	Std	Mean	Median	Std	Mean	Median	Std
Monthly frequency ($T = 744$ observations)									
Market proxy (R_m)	47.8	45.3	47.5	0.29	0.05	0.32	–	–	–
Oracle (R_w)	26.0	14.8	21.9	1.12	1.46	0.68	–	–	–
Joint (γ, ψ, λ)	6.6	6.4	5.5	1.34	1.35	0.40	0.40	0.33	0.20
Annual frequency ($T = 62$ observations)									
Market proxy (R_m)	34.9	0.3	54.4	0.37	0.39	0.33	–	–	–
Oracle (R_w)	33.5	7.9	45.9	0.63	0.95	0.50	–	–	–
Joint (γ, ψ, λ)	10.3	0.0	35.2	0.73	0.72	0.29	0.55	0.65	0.41

Table V: **Short-Sample Power Analysis**
Estimation with 62 years of simulated BKY data (1,000 Monte Carlo paths). True parameters: $\gamma = 10$, $\psi = 1.5$, $\lambda = 0.30$. One-step GMM with identity weighting and 5 BKY dividend claims. Bold: $\hat{\lambda}_{\text{med}} = 0.33$, close to true $\lambda = 0.30$.

	$\hat{\gamma}_{BC}$			$\hat{\psi}_{BC}$		
λ	Median	[5th	95th]	Median	[5th	95th]
0.10	8.3	0.9	9.1	0.437	0.408	6.172
0.15	11.9	0.9	13.1	0.291	0.272	4.115
0.20	15.5	0.8	17.2	0.218	0.204	3.086
0.25	19.2	0.8	21.2	0.175	0.163	2.469
0.30	22.8	0.7	25.3	0.146	0.136	2.057
0.35	26.4	0.7	29.3	0.125	0.117	1.764
0.40	30.1	0.6	33.4	0.109	0.102	1.543
0.48	35.9	0.5	39.8	0.091	0.085	1.286
0.50	37.3	0.5	41.4	0.087	0.082	1.234
0.60	44.6	0.4	49.5	0.073	0.068	1.029
0.70	51.9	0.3	57.6	0.062	0.058	0.882

Table VI: **Bias Correction Sensitivity to λ**
First-order bias correction $\hat{\gamma}_{BC} = \hat{\lambda}(\hat{\gamma}^* - 1) + 1$ and $\hat{\psi}_{BC} = \hat{\psi}^*/\hat{\lambda}$ applied to 1,000 Monte Carlo paths from the BKY simulation (2,000 years monthly). True $\gamma = 10$, $\psi = 1.5$. Brackets show the [5th, 95th] percentile across paths.

Parameter	**Value**	ε_{γ^*}	ε_{ψ^*}	ρ_B	CDK ϱ_k	ρ_B Rank
ρ	0.9820	-10.02	-67.92	655.967	1.1	1
σ	0.0054	-2.03	+0.25	14.834	48.9	2
μ_c	0.0015	+1.07	+0.42	4.331	1.0	3
ν	0.9800	-2.68	-25.79	0.600	97.8	4
ϕ_x	0.0420	-0.25	-1.44	0.388	1.0	5
$\phi_{d,c}$	3.8000	+0.19	+0.87	0.201	1.0	6
ϕ_d	2.3000	-0.10	-1.04	0.087	3.4	7
$\phi_{d,d}$	5.4000	-0.05	-0.78	0.028	1.0	8
σ_w	0.0000	-0.06	-0.62	0.008	1.0	9
μ_d	0.0015	-0.01	-0.08	0.001	1.0	10
Spearman τ_{BE}						0.636
Aggregate ρ_B						676.337

Table VII: **Bias Dark Matter vs. CDK Efficiency Dark Matter**
Elasticities ε_{γ^*}, ε_{ψ^*} computed via central differences ($\pm 0.01\%$ perturbation) with full equilibrium recomputation. $\rho_B(\eta_k) = S_k' V_\vartheta^{-1} S_k \mathrm{Var}(\eta_k)$ with $\mathrm{Var}(\eta_k)$ calibrated at 10% coefficient of variation (1% for ν). CDK ϱ_k from Table IV, Panel A of Chen, Dou, and Kogan (2024), BKY M1 specification. τ_{BE}: Spearman rank correlation between ρ_B and CDK ϱ_k rankings. BKY baseline: $\gamma = 10$, $\psi = 1.5$, $\beta = 0.9989$.

Dark Matter	**Value**	$\partial\gamma^*/\partial\eta$	$\partial\psi^*/\partial\eta$	ρ_B	Bias Red. (%)
π_z	5.12e-04	-13812.88	-1523.8731	100.717	91.0
σ_z^2	3.28e-05	182030.20	20082.0545	286.554	-60.7
σ_{cz}	1.90e-05	-256178.19	-28262.2578	47.800	53.1
σ_{zf}	6.42e-07	78607.43	2442.7817	0.002	0.1

Table VIII: **Bias Dark Matter for Non-Tradable Return Moments**
Sensitivity of pseudo-true preferences to non-tradable return moments $\eta = (\pi_z, \sigma_z^2, \sigma_{cz}, \sigma_{zf})$. "Bias Red." is the percentage reduction in $\|(\gamma, \psi) - (\gamma_0, \psi_0)\|$ from adding each correction individually to the no-correction baseline.

Correction	γ	ψ	err γ (%)	err ψ (%)	Bias Red. (%)
No correction	15.63	2.3907	56.3	59.4	0.0
π_z	10.51	1.4742	5.1	1.7	91.0
σ_z^2	18.63	4.5765	86.3	205.1	-60.7
σ_{cz}	12.66	1.7270	26.6	15.1	53.1
σ_{zf}	15.62	2.3886	56.2	59.2	0.1
All corrections	10.00	1.5000	0.0	0.0	100.0

Table IX: **Dark Matter Channel Decomposition**
Bias decomposition by dark matter channel. "No correction": baseline with $\pi_z = \sigma_z^2 = \sigma_{cz} = \sigma_{zf} = 0$. Each row adds one dark matter moment. "All corrections": exact map recovering true parameters. Bias Red.: percentage reduction in $\|(\gamma, \psi) - (\gamma_0, \psi_0)\|$ relative to no correction.

γ	ψ	γ^*	ψ^*	τ_{BE}	Top Param
5	0.50	106.7	0.0862	0.176	ρ
5	1.01	111.0	0.0694	0.224	ρ
5	1.50	111.1	0.0620	0.224	ρ
5	2.00	111.0	0.0581	0.224	ρ
10	0.50	98.1	0.0697	0.636	ρ
10	1.01	102.2	0.0479	0.636	ρ
10	1.50	102.7	0.0409	0.636	ρ
10	2.00	102.9	0.0374	0.636	ρ
15	1.01	97.3	0.0233	0.673	ρ
15	1.50	98.5	0.0183	0.673	ρ
15	2.00	98.9	0.0158	0.673	ρ
20	1.01	99.5	-0.0021	0.673	ρ
20	1.50	99.6	-0.0040	0.697	σ
20	2.00	99.6	-0.0052	0.612	ρ

Table X: **Bias-Efficiency Concordance Across Calibrations**
τ_{BE}: Spearman rank correlation between ρ_B and CDK ϱ_k rankings at each (γ, ψ) grid point. "Top Param": structural parameter with the largest bias dark matter measure ρ_B. All computations use $\beta = 0.9989$ and BKY parameters from Bansal, Kiku, and Yaron (2007).

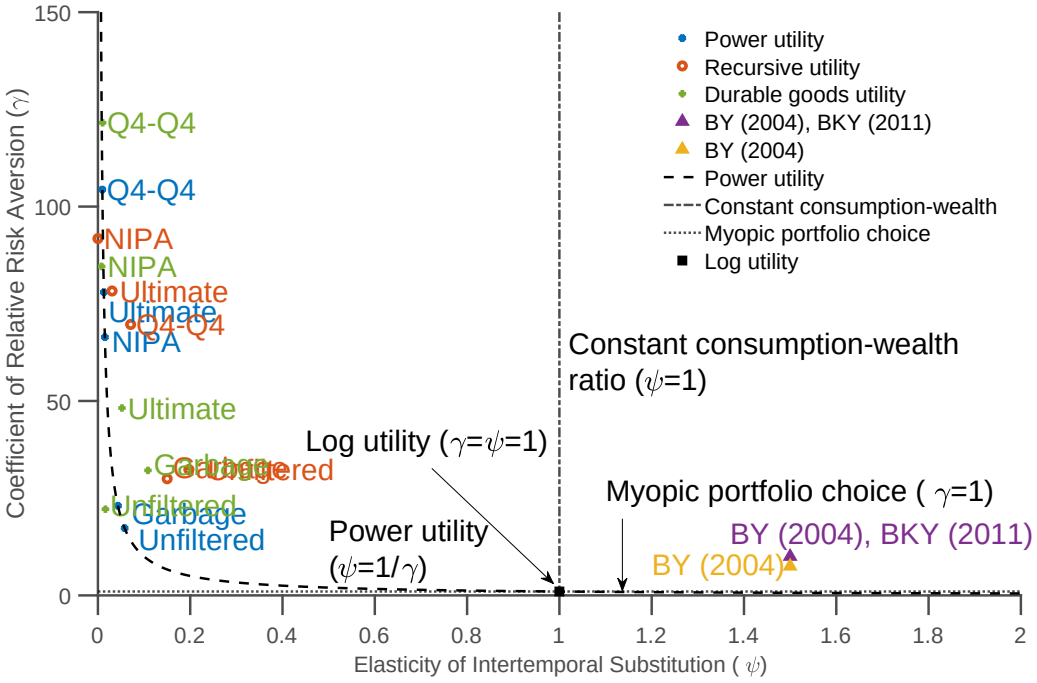
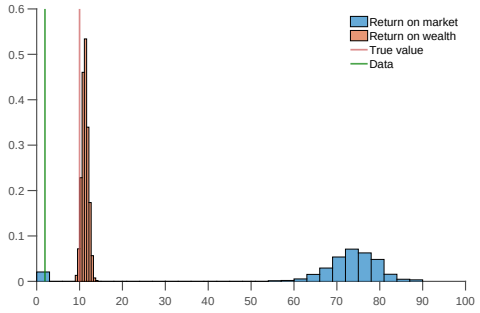
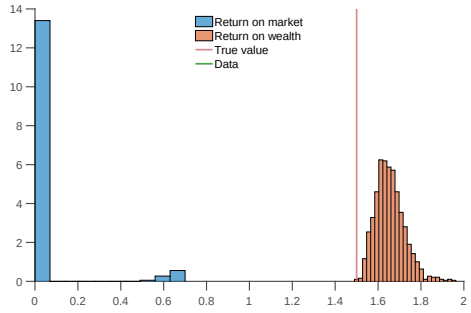


Figure 1: **Non-linear Models: Estimation Summary**

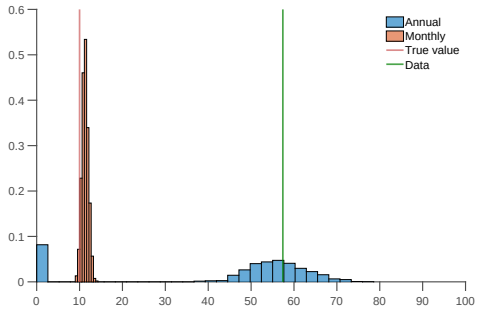
Figure maps the estimated preference parameter space for the Epstein and Zin (1989) recursive utility framework. Each point represents GMM estimates of the coefficient of relative risk aversion (γ, y-axis) and the elasticity of intertemporal substitution (ψ, x-axis) from non-linear estimation using different consumption measures (NIPA, Ultimate, Q4-Q4, Unfiltered, and Garbage) and preference specifications (power utility shown as circles, recursive utility as squares, and durable goods utility as triangles). Test assets include the market excess return, risk-free rate, and 25 Fama-French size and book-to-market portfolios. The diagonal line $\gamma = 1/\psi$ represents power utility. The vertical line at $\psi = 1$ indicates a constant consumption-wealth ratio, while the horizontal line at $\gamma = 1$ represents myopic portfolio choice. Calibrated values from Bansal and Yaron (2004) (BY 2004) and Bansal, Kiku, and Yaron (2007) (BKY 2011) are shown for comparison. Sample: 1961–2022.



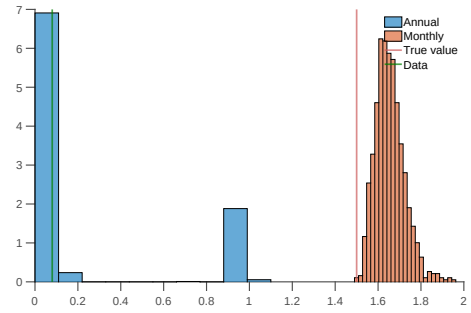
(a) Return on wealth measure: RRA



(b) Return on wealth measure: EIS

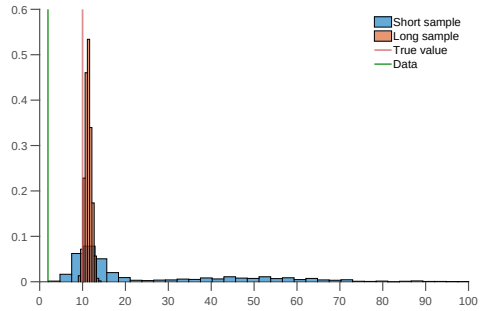


(c) Temporal aggregation: RRA

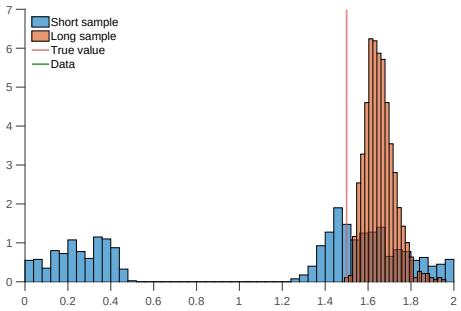


(d) Temporal aggregation: EIS

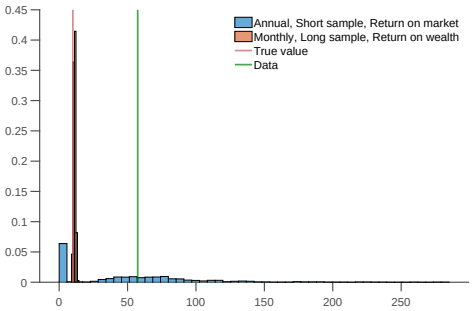
Figure 2: **Distribution of Estimated Parameters: Bias Decomposition (Part I)**
Histograms of estimated preference parameters across Monte Carlo paths in the BKY long-run risk model. True parameters: $\gamma = 10$, $\psi = 1.5$. Red vertical line: true value. Green vertical line: empirical estimate from NIPA data. Panels (a)–(b): Monthly long sample comparing estimation using the return on market (R_m) versus the return on wealth (R_w). Panels (c)–(d): Long sample using R_w, comparing annual versus monthly frequency.



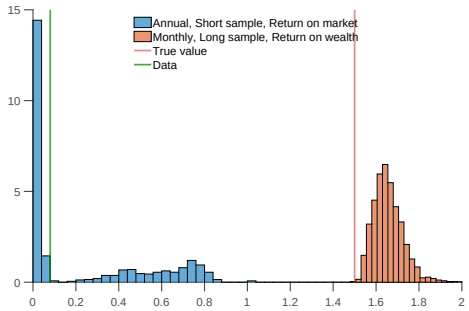
(a) Small sample: RRA



(b) Small sample: EIS



(c) All biases: RRA



(d) All biases: EIS

Figure 3: **Distribution of Estimated Parameters: Bias Decomposition (Part II)**
Histograms of estimated preference parameters across Monte Carlo paths in the BKY long-run risk model. True parameters: $\gamma = 10$, $\psi = 1.5$. Red vertical line: true value. Green vertical line: empirical estimate from NIPA data. Panels (a)–(b): Monthly data using R_w, comparing short sample ($T = 780$ months) versus long sample ($T = 24{,}000$ months). Panels (c)–(d): Worst case combining all three biases (annual frequency, short sample, return on market) versus the best case (monthly, long sample, return on wealth).

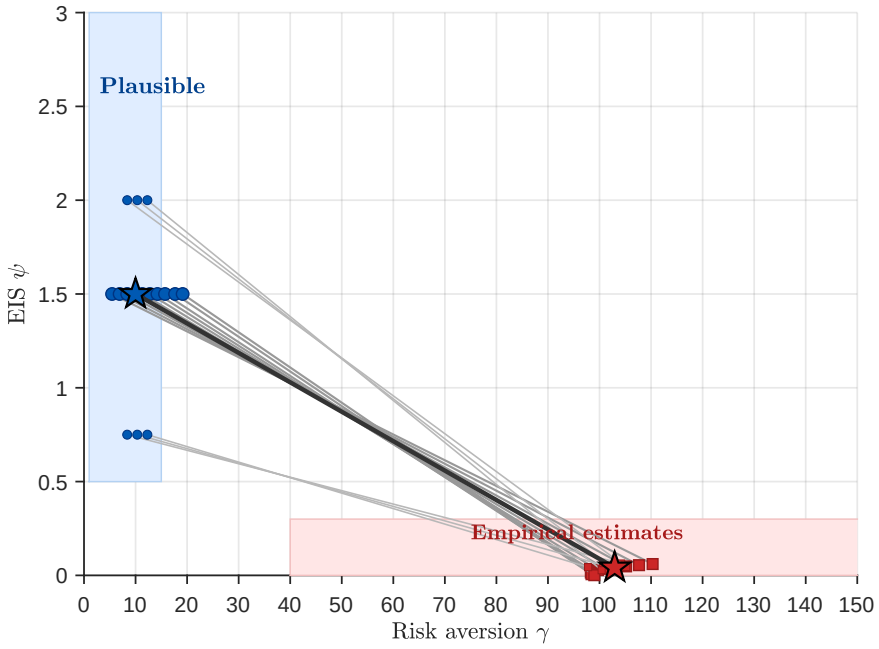
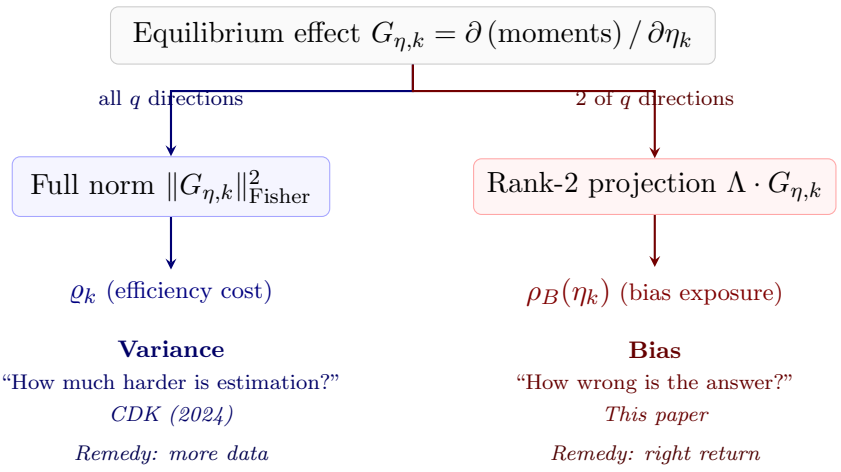
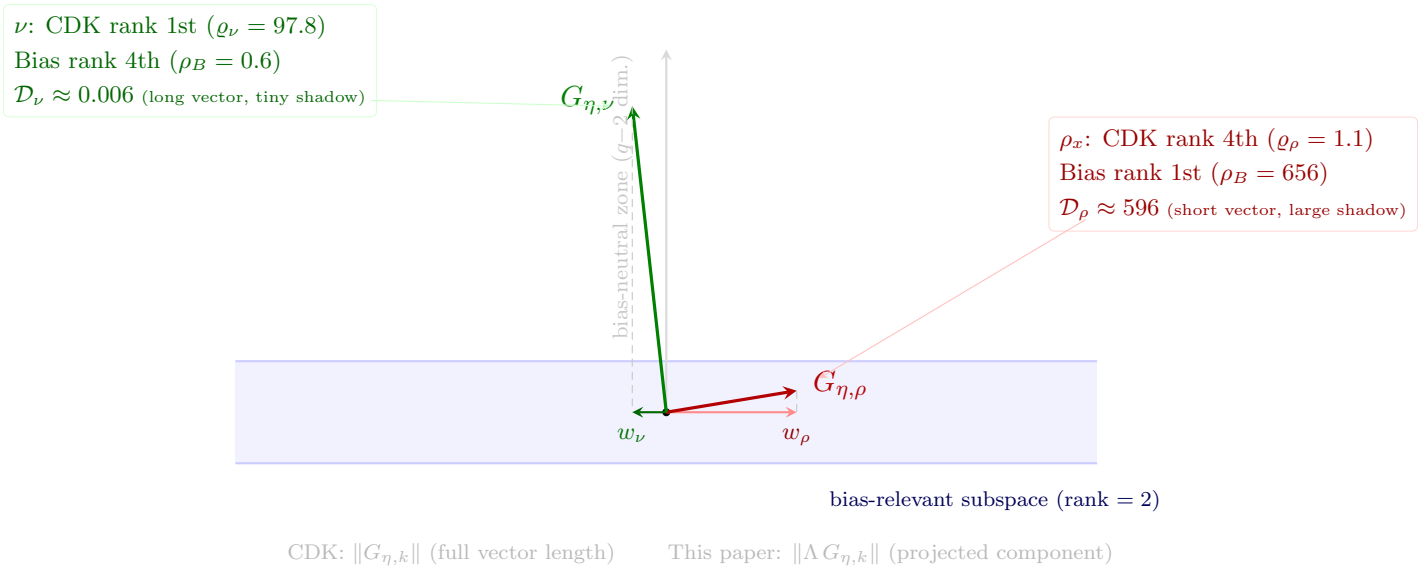


Figure 4: **Mapping from True to Pseudo-True Parameters**

The figure illustrates how wealth mismeasurement distorts preference inference. Each arrow connects a true (γ, ψ) pair (blue circles) to its full Euler pseudo-true (γ^*, ψ^*) (red squares), computed from Proposition 4. True parameters span the economically plausible region (blue shading: $\gamma \in [1, 15]$, $\psi \in [0.5, 3]$) for three values of the EIS ($\psi = 0.75, 1.5, 2.0$). The starred markers highlight the BKY benchmark ($\gamma = 10$, $\psi = 1.5$). Despite wide variation in true preferences, all pseudo-true estimates cluster in a narrow pathological band (red shading: $\gamma^* \in [40, 200]$, $\psi^* \approx 0$) –precisely the region where four decades of empirical work has placed them. The convergence of arrows to a single region is the visual counterpart of the near-invariance result: empirical estimates of risk aversion are essentially uninformative about the true parameter.



(a) Same equilibrium Jacobian, different lens



(b) Why rankings diverge: projection geometry in moment space

Figure 5: **Same Jacobian, Different Lens: CDK and Bias**
Panel (a): Both analyses start from the equilibrium Jacobian $G_{\eta,k}$, which measures how structural parameter η_k shifts observable moments through the model's equilibrium. Chen, Dou, and Kogan (2024) evaluate the full norm across all q moment directions, yielding the efficiency cost ϱ_k that governs estimation variance. We project through the rank-two influence function Λ, yielding the bias exposure $\rho_B(\eta_k)$ that governs how much the pseudo-true preference estimates shift. Panel (b): Each parameter's equilibrium effect $G_{\eta,k}$ decomposes into a bias-relevant component w_k (horizontal, survives Λ) and a bias-neutral component (vertical, annihilated by Λ). Volatility persistence ν (green): long vector, tiny projection ($\mathcal{D}_\nu \approx 0.006$). Long-run risk persistence ρ (red): shorter vector, but nearly all of it falls in the bias-relevant subspace ($\mathcal{D}_\rho \approx 596$). The ranking reversal follows from the rank-two structure of Λ.

REFERENCES

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro, 2017, Measuring the sensitivity of parameter estimates to estimation moments, *Quarterly Journal of Economics* 132, 1553–1592.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2007, Risks for the long run: Estimation and inference.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *The Journal of Finance* 59, 1481–1509.

Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Journal of Economics* 121, 823–866.

Breeden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.

Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.

Campbell, John Y., 2003, Consumption-based asset pricing, *Handbook of the Economics of Finance* 1, 893–887.

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

Chen, Hui, Winston Wei Dou, and Leonid Kogan, 2024, Measuring "dark matter" in asset pricing models, *The Journal of Finance* 79, 843–902.

Chen, Xiaohong, Jack Favilukis, and Sydney C. Ludvigson, 2013, An estimation of economic models with recursive preferences: Economic models with recursive preferences, *Quantitative Economics* 4, 39–83.

Cochrane, John H., 2005, *Asset Pricing (Revised Edition)* (Princeton University Press).

Constantinides, George M, 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of political Economy* 98, 519–543.

Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Epstein, Larry G., and Stanley E. Zin, 1991, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis, *Journal of Political Economy* 99, 263–286.

Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *The Journal of Finance* 47, 427–465.

Favilukis, Jack, and Xiaoji Lin, 2016, Wage rigidity: A quantitative solution to several asset pricing puzzles, *Review of Financial Studies* 29, 523–564.

Ghosh, Anisha, Christian Julliard, and Alex P. Taylor, 2017, What is the consumption-capm missing? an information-theoretic framework for the analysis of asset pricing models, *The Review of Financial Studies* 30, 442–504.

Guo, Hui, and Robert Savickas, 2019, Understanding the risk-return relation: The aggregate wealth proxy actually matters, *Journal of Business & Economic Statistics* 37, 631–642.

Hall, Alastair R., and Atsushi Inoue, 2003, The large sample behaviour of the generalized method of moments estimator in misspecified models, *Journal of Econometrics* 114, 361–394.

Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 339–357.

Hansen, Lars Peter, and Ravi Jagannathan, 1997, Assessing specification errors in stochastic discount factor models, *The Journal of Finance* 52, 557–590.

Hansen, Lars Peter, and Kenneth J. Singleton, 1982, Generalized instrumental variables estimation of nonlinear rational expectations models, *Econometrica* 50, 1269–1286.

Hansen, Lars Peter, and Kenneth J. Singleton, 1983, Stochastic consumption, risk aversion, and the temporal behavior of asset returns, *Journal of Political Economy* 91, 249–265.

Jagannathan, Ravi, and Yong Wang, 2007, Lazy investors, discretionary consumption, and the cross-section of stock returns, *The Journal of Finance* 62, 1623–1661.

Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional capm and the cross-section of expected returns, *The Journal of Finance* 51, 3–53.

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor, 2019, The rate of return on everything, 1870–2015, *Quarterly Journal of Economics* 134, 1225–1298.

Kleibergen, Frank, 2009, Tests of risk premia in linear factor models, *Journal of Econometrics* 149, 149–173.

Kreps, David M, and Evan L Porteus, 1978, Temporal resolution of uncertainty and dynamic choice theory, *Econometrics* 46, 185–200.

Kroencke, Tim A., 2017, Asset pricing without garbage, *The Journal of Finance* 72, 47–98.

Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *The Journal of Finance* 56, 815–849.

Lucas, Jr., Robert E., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Ludvigson, Sydney C., 2013, Advances in consumption-based asset pricing: Empirical tests, *Handbook of the Economics of Finance* .

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2013, The wealth-consumption ratio, *Review of Asset Pricing Studies* 3, 38–94.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Merton, Robert C, 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Parker, Jonathan A., and Christian Julliard, 2005, Consumption risk and the cross section of expected returns, *Journal of Political Economy* 113, 185–222.

Rietz, Thomas A., 1988, The equity premium: A solution, *Journal of Monetary Economics* 22, 117–131.

Roll, Richard, 1977, A critique of the asset pricing theory's tests part i: On past and potential testability of the theory, *Journal of financial economics* 4, 129–176.

Santos, Tano, and Pietro Veronesi, 2006, Labor income and predictable stock returns, *Review of Financial Studies* 19, 1–44.

Savov, Alexi, 2011, Asset pricing with garbage, *The Journal of Finance* 66, 177–201.

Shanken, Jay, 1987, Multivariate proxies and asset pricing relations: Living with the Roll critique, *Journal of Financial Economics* 18, 91–110.

Stock, James H, and Jonathan H Wright, 2000, Gmm with weak identification, *Econometrica* 68, 1055–1096.

Weber, Christian E., 2000, "rule-of-thumb" consumption, intertemporal substitution, and risk aversion, *Journal of Business and Economic Statistics* 18, 497–502.

Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401–421.

Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *The Quarterly Journal of Economics* 105, 29–42.

White, Halbert, 1982, Maximum likelihood estimation of misspecified models, *Econometrica* 50, 1–25.

Yogo, Motohiro, 2006, A consumption-based explanation of expected stock returns, *The Journal of Finance* 61, 539–580.

Appendix A. Proofs of Main Results

A. Proof of Theorem 2: Variance-Matching Bias Characterization

The true log SDF in an Epstein–Zin economy is

$$m_{t+1} = \theta \log \beta - \frac{\theta}{\psi}\Delta c_{t+1} + (\theta - 1) r_{w,t+1}, \tag{A1}$$

where $\theta = (1-\gamma)/(1-1/\psi)$. Total wealth decomposes as $r_w = \lambda r_m + (1-\lambda) r_z$, where r_m is the observed market return, r_z is unobserved, and $\mathrm{Cov}(r_m, r_z) = \rho_{mz}\sigma_m\sigma_z$. The econometrician, not observing r_z, uses only r_m as a proxy:

$$m^*_{t+1} = \theta^* \log \beta - \frac{\theta^*}{\psi}\Delta c_{t+1} + (\theta^* - 1) r_{m,t+1}. \tag{A2}$$

Focusing on the wealth component of the log SDF, the variance of the true log SDF is

$$\mathrm{Var}(m) \approx (\theta-1)^2 \mathrm{Var}(r_w) = (\theta-1)^2[\lambda^2\sigma_m^2 + (1-\lambda)^2\sigma_z^2 + 2\lambda(1-\lambda)\rho_{mz}\sigma_m\sigma_z], \tag{A3}$$

while the variance of the misspecified SDF is $\mathrm{Var}(m^*) \approx (\theta^* - 1)^2\sigma_m^2$. The pseudo-true parameter θ^* equates these:

$$(\theta^* - 1)^2\sigma_m^2 = (\theta-1)^2[\lambda^2\sigma_m^2 + (1-\lambda)^2\sigma_z^2 + 2\lambda(1-\lambda)\rho_{mz}\sigma_m\sigma_z]. \tag{A4}$$

Dividing both sides by σ_m^2 and defining $\tilde{\eta} = (1-\lambda)^2\sigma_z^2/\sigma_m^2 + 2\lambda(1-\lambda)\rho_{mz}\sigma_z/\sigma_m$ gives $(\theta^* - 1)^2 = (\theta - 1)^2(\lambda^2 + \tilde{\eta})$. Taking the positive square root (required for continuity as $\lambda \to 1$):

$$\theta^* = 1 + (\theta - 1)\sqrt{\lambda^2 + \tilde{\eta}}. \tag{A5}$$

Converting to γ^* via $\gamma = 1 - \theta(1 - 1/\psi)$ and the identity $(\theta - 1)(1 - 1/\psi) = 1/\psi - \gamma$:

$$\begin{aligned}\gamma^* &= 1 - \theta^*(1 - 1/\psi) = 1 - [1 + (\theta - 1)\sqrt{\lambda^2 + \tilde{\eta}}](1 - 1/\psi) \\ &= \frac{1}{\psi} - (\theta - 1)(1 - 1/\psi)\sqrt{\lambda^2 + \tilde{\eta}} = \frac{1}{\psi} + (\gamma - 1/\psi)\sqrt{\lambda^2 + \tilde{\eta}}.\end{aligned} \tag{A6}$$

Substituting $\tilde{\eta} = (1-\lambda)^2\sigma_z^2/\sigma_m^2 + 2\lambda(1-\lambda)\rho_{mz}\sigma_z/\sigma_m$ completes the proof. When $\rho_{mz} = 0$, this reduces to the orthogonal case with $\tilde{\eta} = (1-\lambda)^2\sigma_z^2/\sigma_m^2$. □

B. Proof of Proposition 8: Eta Expansion

Expanding $\sqrt{\lambda^2 + \tilde{\eta}}$ around $\tilde{\eta} = 0$:

$$\sqrt{\lambda^2 + \tilde{\eta}} = \lambda\sqrt{1 + \tilde{\eta}/\lambda^2} = \lambda + \frac{\tilde{\eta}}{2\lambda} - \frac{\tilde{\eta}^2}{8\lambda^3} + O(\tilde{\eta}^3). \tag{A7}$$

Substituting into $\gamma^* = 1/\psi + (\gamma - 1/\psi)\sqrt{\lambda^2 + \tilde{\eta}}$:

$$\begin{aligned}\gamma^* &= \frac{1}{\psi} + (\gamma - 1/\psi)\left[\lambda + \frac{\tilde{\eta}}{2\lambda} - \frac{\tilde{\eta}^2}{8\lambda^3}\right] + O(\tilde{\eta}^3) \\ &= \lambda\gamma + (1-\lambda)\frac{1}{\psi} + \frac{\gamma - 1/\psi}{2\lambda}\tilde{\eta} - \frac{\gamma - 1/\psi}{8\lambda^3}\tilde{\eta}^2 + O(\tilde{\eta}^3),\end{aligned} \tag{A8}$$

where $\tilde{\eta} = (1-\lambda)^2\sigma_z^2/\sigma_m^2 + 2\lambda(1-\lambda)\rho_{mz}\sigma_z/\sigma_m$. Reading off coefficients gives $\gamma_0^* = \lambda\gamma + (1-\lambda)/\psi$, $\gamma_1^* = (\gamma - 1/\psi)/(2\lambda)$, and $\gamma_2^* = -(\gamma - 1/\psi)/(4\lambda^3)$. □

C. Proof of Theorem 3: EIS Bias

The EIS bias has two conceptually distinct sources: variance matching (Theorem 2) and risk premium matching. We treat each in turn.

Variance matching alone implies $\psi^* = \psi$. From $\theta = (1-\gamma)/(1-1/\psi)$, the identity $\theta - 1 + \gamma = \theta/\psi$ holds. The inverse EIS satisfies $1/\psi = (\theta - 1 + \gamma)/\theta$, and likewise for the pseudo-true. Substituting the Theorem 2 expressions $\theta^* = 1 + (\theta - 1)\rho$ and $\gamma^* = 1/\psi + (\gamma - 1/\psi)\rho$:

$$\begin{aligned}\frac{1}{\psi^*} &= \frac{\theta^* - 1 + \gamma^*}{\theta^*} = \frac{(\theta-1)\rho + 1/\psi + (\gamma - 1/\psi)\rho}{1 + (\theta-1)\rho} = \frac{1/\psi + \rho[(\theta-1) + \gamma - 1/\psi]}{1 + (\theta-1)\rho} \\ &= \frac{1/\psi + \rho(\theta-1)/\psi}{1 + (\theta-1)\rho} = \frac{1}{\psi},\end{aligned} \tag{A9}$$

where we used $(\theta - 1) + \gamma - 1/\psi = (\theta - 1)/\psi$. Thus the SDF variance channel does not distort the EIS.

Risk premium matching, however, does distort ψ^*. Under log-normality, the risk premium for asset j is $-\text{Cov}(m, r_j)$. Writing $a = \theta/\psi$ and $b = \theta - 1$ for the consumption and wealth

loadings, the pseudo-true must satisfy

$$-(a^* - a)\sigma_{cj} + (b^* - b\lambda)\sigma_{mj} = b(1-\lambda)\sigma_{zj}, \tag{A10}$$

where $\sigma_{cj} = \text{Cov}(\Delta c, r_j)$, $\sigma_{mj} = \text{Cov}(r_m, r_j)$, and $\sigma_{zj} = \text{Cov}(r_z, r_j)$. When $\sigma_{zj} \neq 0$, the misspecified model lacks the r_z factor and must absorb the unobserved risk premium through Δc and r_m, distorting $a^* = \theta^*/\psi^*$ and hence the pseudo-true EIS.

Expanding around $\lambda = 1$ with $\varepsilon = 1 - \lambda$, the $O(\varepsilon)$ correction from the market risk premium condition ($j = m$) yields

$$a_1\sigma_{cm} - \theta_1\sigma_m^2 = b(\sigma_m^2 - \rho_{mz}\sigma_m\sigma_z), \tag{A11}$$

where $a^* = a + \varepsilon a_1$ and $\theta^* = \theta + \varepsilon\theta_1$. A second moment condition (from a cross-sectional asset) provides a second equation. Solving this 2×2 system determines (a_1, θ_1), and hence $1/\psi^* = a^*/\theta^*$ to first order. The formula in the theorem statement follows from evaluating this system under the approximation $b^* \approx b\lambda$. □

D. Proof of Theorem 9: Identification Bounds

From the zeroth-order formula $\gamma_0^* = \lambda\gamma + (1-\lambda)/\psi$, when $\gamma \gg 1/\psi$ we have $\gamma^* - 1 \approx \lambda(\gamma - 1)$, so $\gamma = \lambda(\gamma^* - 1) + 1$. If $\lambda \in [\underline{\lambda}, \bar{\lambda}]$, then

$$\gamma \in [\underline{\lambda}(\gamma^* - 1) + 1,\ \bar{\lambda}(\gamma^* - 1) + 1]. \tag{A12}$$

These bounds are sharp because the mapping $\gamma(\lambda) = \lambda(\gamma^* - 1) + 1$ is monotonically increasing in λ (since $\gamma^* > 1$), and for each $\lambda \in [\underline{\lambda}, \bar{\lambda}]$ there exists a data-generating process consistent with the observed γ^*. □

E. *Proof of Theorem 10: Identification Strength*

The identification strength measure is

$$\mu_N = \sqrt{N} \cdot \frac{\|\nabla_\theta g(\theta_0)\|_W}{\sqrt{\text{Var}(\sqrt{N}\bar{g}_N(\theta_0))}}. \tag{A13}$$

When the econometrician uses r_m instead of r_w, the Jacobian $\nabla_\theta \tilde{g}$ is attenuated because r_m provides only a fraction λ of the variation in r_w: $\nabla_\theta \tilde{g} \approx \lambda \cdot \nabla_\theta g + O((1-\lambda)^2)$. The variance of the sample moments is simultaneously inflated by the omitted variable. The ratio of attenuated numerator to inflated denominator gives

$$\mu_{N,proxy} \leq \lambda \cdot \mu_{N,true} \cdot \left(1 - \frac{(1-\lambda)^2 \sigma_z^2}{2\sigma_m^2}\right) + O((1-\lambda)^3). \tag{A14}$$

□

F. *Proof of Proposition 1: Identification of λ*

It suffices to verify that the $M \times 3$ Jacobian matrix $D = [D_\gamma, D_\psi, D_\lambda]$, where $D_\gamma = \partial g / \partial \gamma$, etc., has rank 3 at the true parameter values. The λ-column is

$$\frac{\partial g_j}{\partial \lambda} = \mathbb{E}\left[\frac{\partial M_{t+1}}{\partial \lambda} R_{j,t+1}\right], \tag{A15}$$

where

$$\frac{\partial M_{t+1}}{\partial \lambda} = (\theta - 1) M_{t+1} \frac{R_{m,t+1} - R_{h,t+1}}{\hat{R}_{w,t+1}}. \tag{A16}$$

This derivative is non-zero whenever $R_m \neq R_h$ (condition (i)) and $\theta \neq 1$ (condition (ii)). Because the λ-derivative depends on the level return differential $R_m - R_h$ –functionally independent of the log-level terms Δc and $\log \hat{R}_w$ driving the γ- and ψ-derivatives –the three Jacobian columns span distinct directions in the moment space when $N \geq 3$ test assets have heterogeneous exposures to market and human capital risk. □

Appendix B. Special Cases

In the following, $\rho = \sqrt{\lambda^2 + (1-\lambda)^2\sigma_z^2/\sigma_m^2 + 2\lambda(1-\lambda)\rho_{mz}\sigma_z/\sigma_m}$ is the scaling factor from Theorem 2.

A. Full Observation ($\lambda \to 1$)

As $\lambda \to 1$: $\tilde{\eta} \to 0$, $\rho \to 1$, and $\gamma^* \to 1/\psi + (\gamma - 1/\psi) = \gamma$. Full observation of wealth leads to correct identification.

B. No Observation ($\lambda \to 0$)

As $\lambda \to 0$: $\rho \to \sigma_z/\sigma_m$, and $\gamma^* \to 1/\psi + (\gamma - 1/\psi)\sigma_z/\sigma_m$. The bias depends entirely on the relative volatility of non-traded and traded returns.

C. No Unobserved Variance ($\sigma_z^2 \to 0$)

As $\sigma_z^2 \to 0$: $\tilde{\eta} \to 0$, $\rho \to \lambda$, and $\gamma^* \to \lambda\gamma + (1-\lambda)/\psi = \gamma_0^*$. The zeroth-order approximation becomes exact.

D. Orthogonal Components ($\rho_{mz} = 0$)

When $\rho_{mz} = 0$: $\tilde{\eta} = (1-\lambda)^2\sigma_z^2/\sigma_m^2 = \eta$, recovering the simpler expression $\rho = \sqrt{\lambda^2 + \eta}$.

E. Time-Additive Expected Utility ($\gamma = 1/\psi$)

When $\gamma = 1/\psi$: $\theta = 1$ and $\gamma^* = 1/\psi + 0 \cdot \rho = 1/\psi = \gamma$. There is no bias regardless of λ, σ_z^2/σ_m^2, or ρ_{mz}, because the SDF does not depend on the wealth return.

F. Exact Identification ($\rho = 1$)

Setting $\rho = 1$ requires $\lambda^2 + (1-\lambda)^2\sigma_z^2/\sigma_m^2 + 2\lambda(1-\lambda)\rho_{mz}\sigma_z/\sigma_m = 1$. When this condition holds, $\gamma^* = \gamma$ exactly, even though only fraction λ of wealth is observed. This is a knife-edge: unobserved wealth variance (modulated by ρ_{mz}) exactly compensates for partial observation.

Appendix C. Additional Theoretical Results

A. First-Order Bias and Zeroth-Order Approximation

Theorem 6 (First-Order Bias)**:** *Under Assumptions 1 and 2, the pseudo-true aggregation parameter satisfies:*

$$\theta^* = 1 + (\theta - 1)\lambda + O(\tilde{\eta}). \tag{C1}$$

Corollary 7 (Zeroth-Order Risk Aversion Bias)**:** *The zeroth-order pseudo-true risk aversion is:*

$$\gamma_0^* = \lambda\gamma + (1 - \lambda)\frac{1}{\psi} \tag{C2}$$

B. Taylor Expansion of the Exact Solution

Proposition 8 (Eta Expansion)**:** *The pseudo-true risk aversion expands as:*

$$\gamma^* = \underbrace{\lambda\gamma + (1 - \lambda)\frac{1}{\psi}}_{\gamma_0^*} + \underbrace{\frac{\gamma - 1/\psi}{2\lambda}}_{\gamma_1^*}\tilde{\eta} - \underbrace{\frac{\gamma - 1/\psi}{8\lambda^3}}_{\gamma_2^*/2}\tilde{\eta}^2 + O(\tilde{\eta}^3). \tag{C3}$$

C. Identification Bounds

Theorem 9 (Identification Bounds)**:** *Without knowledge of λ, the identified set for γ given observed γ^* is:*

$$\mathcal{G}(\gamma^*) = \left\{\gamma : \gamma \in \left[\underline{\lambda}(\gamma^* - 1) + 1,\ \bar{\lambda}(\gamma^* - 1) + 1\right]\right\}, \tag{C4}$$

where $[\underline{\lambda}, \bar{\lambda}]$ represents prior bounds on the wealth share.

For example, with $\gamma^* = 70$: prior $\lambda \in [0.2, 0.5]$ gives $\gamma \in [14.8, 35.5]$; prior $\lambda \in [0.3, 0.4]$ gives $\gamma \in [21.7, 28.6]$.

D. Identification Strength

Theorem 10 (Identification Strength)**:** *Define the identification strength measure* $\mu_N = \sqrt{N} \cdot \|\nabla_\theta g(\theta_0)\|_W / \sqrt{Var(\sqrt{N}\bar{g}_N(\theta_0))}$*. Then:*

$$\mu_{N,proxy} \leq \lambda \cdot \mu_{N,true} \cdot \left(1 - \frac{(1-\lambda)^2 \sigma_z^2}{2\sigma_m^2}\right) + O((1-\lambda)^3). \tag{C5}$$

Identification strength is reduced by a factor of approximately λ. For $\lambda \approx 0.3$, strength is reduced by roughly 70 percent, explaining why consumption-based models appear weakly identified (Stock and Wright, 2000; Kleibergen, 2009).

Appendix D. Additional Bias Correction Details

A. Second-Order Iterative Correction

Initialize with $\gamma^{(0)} = \hat{\lambda}(\hat{\gamma}^* - 1) + 1$. At each iteration k, compute

$$\rho^{(k)} = \sqrt{\hat{\lambda}^2 + (1-\hat{\lambda})^2 \, \hat{\sigma}_z^2/\hat{\sigma}_m^2 + 2\hat{\lambda}(1-\hat{\lambda}) \, \hat{\rho}_{mz} \, \hat{\sigma}_z/\hat{\sigma}_m} \tag{D1}$$

and update $\gamma^{(k+1)} = 1/\hat{\psi}^{(k)} + (\hat{\gamma}^* - 1/\hat{\psi}^{(k)})/\rho^{(k)}$. Terminate when $|\gamma^{(k+1)} - \gamma^{(k)}|$ falls below a specified tolerance. Convergence is achieved within 3–5 iterations for $\lambda > 0.2$.

B. General EIS Correction Formula

The full EIS correction formula is

$$\hat{\psi}_{BC} = \frac{1}{\frac{1/\hat{\psi}^* - (1-\hat{\lambda})\hat{\gamma}_{BC}\hat{\rho}_{cz}\hat{\sigma}_z/\hat{\sigma}_c}{\hat{\lambda}\left(1 + \frac{(1-\hat{\lambda})\hat{\gamma}_{BC}\hat{\rho}_{cm}\hat{\sigma}_c}{\hat{\lambda}\hat{\sigma}_m}\right)}}. \tag{D2}$$

When $\hat{\rho}_{cz}$ and $\hat{\rho}_{cm}$ are unknown or set to zero, this simplifies to $\hat{\psi}_{BC} = \hat{\psi}^*/\hat{\lambda}$.

C. Alternative Methods for Estimating λ

Direct measurement from Flow of Funds data yields estimates ranging from $\lambda \approx 0.3$ to $\lambda \approx 0.5$, depending on the treatment of housing wealth and the valuation of human capital. Alternatively, if one has priors on the true γ, inversion gives $\hat{\lambda} \approx (\gamma_{prior} - 1)/(\hat{\gamma}^* - 1)$.

Appendix E. Proofs of Dark Matter Results

This appendix proves the formal results stated in Section IV.

A. Chain Rule Decomposition

Proposition 11 (Chain Rule Decomposition)**:** *Let $g(\vartheta; \eta) : \mathbb{R}^2 \times \mathbb{R}^{d_\eta} \to \mathbb{R}^q$ collect q Euler equation residuals, where $\vartheta = (\gamma, \psi)$ and η parameterises the non-tradable return moments. Let $\vartheta^*(\eta)$ denote the pseudo-true parameters that solve the population GMM problem. Suppose the Jacobian $G = \partial g/\partial \vartheta|_{\vartheta^*}$ has full column rank. Then the bias sensitivity matrix $S = \partial \vartheta^*/\partial \eta$ satisfies*

$$S = \Lambda \cdot G_\eta, \tag{E1}$$

where $\Lambda = -(G'WG)^{-1}G'W \in \mathbb{R}^{2\times q}$ is the GMM influence function and $G_\eta = \partial g/\partial \eta|_{(\vartheta^*,\eta)} \in \mathbb{R}^{q\times d_\eta}$ is the equilibrium response Jacobian. The decomposition is exact when $q = 2$ (just-identified) and holds to first order in $\|g(\vartheta^*; \eta)\|$ when $q > 2$ (overidentified).

Proof. **Just-identified case** ($q = 2$). The pseudo-true $\vartheta^*(\eta)$ solves $g(\vartheta^*(\eta), \eta) = 0$ identically in η. Differentiating both sides with respect to η:

$$\frac{\partial g}{\partial \vartheta}\frac{\partial \vartheta^*}{\partial \eta} + \frac{\partial g}{\partial \eta} = 0 \quad \Longrightarrow \quad G \cdot S = -G_\eta. \tag{E2}$$

Since G is 2×2 and nonsingular by assumption, $S = -G^{-1}\, G_\eta$. When $q = 2$, the influence function reduces to $\Lambda = -(G'WG)^{-1}G'W = -G^{-1}$ for any nonsingular W, so $S = \Lambda\, G_\eta$.

Overidentified case ($q > 2$). The pseudo-true minimises the population GMM criterion

$Q(\vartheta;\eta) = g(\vartheta;\eta)'\, W\, g(\vartheta;\eta)$. The first-order condition is

$$F(\vartheta,\eta) \;\equiv\; G(\vartheta)'\, W\, g(\vartheta;\eta) = 0. \tag{E3}$$

By the implicit function theorem:

$$S = -\left[\frac{\partial F}{\partial \vartheta}\right]^{-1} \frac{\partial F}{\partial \eta} = -\big[H + G'WG\big]^{-1} G'W\, G_\eta, \tag{E4}$$

where $H = \sum_{j=1}^{q} [W\, g(\vartheta^*;\eta)]_j \cdot \nabla^2_\vartheta g_j(\vartheta^*)$ collects the Hessian terms weighted by the residual vector $g(\vartheta^*;\eta)$. In the just-identified case $g(\vartheta^*;\eta) = 0$, so $H = 0$ and the formula reduces to $S = -(G'WG)^{-1}G'WG_\eta = \Lambda\, G_\eta$. In the overidentified case $H = O(\|g(\vartheta^*;\eta)\|)$, so $S = \Lambda\, G_\eta + O(\|g(\vartheta^*;\eta)\|)$. □

Remark 3. In the BKY calibration, the core Euler system ($q = 2$: risk-free rate and market risk premium) is just-identified for $\vartheta = (\gamma, \psi)$, and the chain rule holds exactly. When cross-sectional moment conditions are added ($q > 2$), the approximation error is controlled by the norm of the Euler residuals at the pseudo-true, which is small relative to the bias effects documented in Section IV.

B. Rank-Two Property of the Bias-Weighted Norm

Proposition 12 (Rank-Two Property)**:** *Let $\Lambda \in \mathbb{R}^{2\times q}$ with $q \geq 3$ have full row rank, and let $V_\vartheta \in \mathbb{R}^{2\times 2}$ be positive definite. Define $W_B = \Lambda'\, V_\vartheta^{-1}\, \Lambda \in \mathbb{R}^{q\times q}$. Then:*

(i) $\operatorname{rank}(W_B) = 2$.

(ii) $\dim(\ker(W_B)) = q - 2$.

(iii) For any $G_{\eta,k} \in \ker(W_B)$: $\Lambda\, G_{\eta,k} = 0$ and $\rho_B(\eta_k) = 0$.

Proof. (i) Since V_ϑ^{-1} is positive definite, it admits a Cholesky decomposition $V_\vartheta^{-1} = L'L$ with $L \in \mathbb{R}^{2\times 2}$ nonsingular. Then $W_B = (L\Lambda)'\,(L\Lambda)$. Because L is nonsingular, $\operatorname{rank}(L\Lambda) = \operatorname{rank}(\Lambda) = 2$. For any matrix A, $\operatorname{rank}(A'A) = \operatorname{rank}(A)$, so $\operatorname{rank}(W_B) = 2$.

(ii) By the rank-nullity theorem, $\dim(\ker(W_B)) = q - \operatorname{rank}(W_B) = q - 2$.

(iii) If $W_B\, x = 0$, then $x'\, W_B\, x = x'\Lambda'\, V_\vartheta^{-1}\, \Lambda\, x = (\Lambda x)'\, V_\vartheta^{-1}\, (\Lambda x) = 0$. Since V_ϑ^{-1} is positive definite, this implies $\Lambda\, x = 0$. Setting $x = G_{\eta,k}$ gives $S_k = \Lambda\, G_{\eta,k} = 0$, hence $\rho_B(\eta_k) = S_k'\, V_\vartheta^{-1}\, S_k\, \mathrm{Var}(\eta_k) = 0$. □

C. SVD Decomposition and Amplification

Corollary 13 (SVD Amplification)**:** *Let $\Lambda = U\Sigma V'$ be the thin singular value decomposition, where $U \in \mathbb{R}^{2\times 2}$ is orthogonal, $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2)$ with $\sigma_1 \geq \sigma_2 > 0$, and $V \in \mathbb{R}^{q\times 2}$ has orthonormal columns. For any $G_{\eta,k} \in \mathbb{R}^q$:*

(i) $G_{\eta,k}$ decomposes orthogonally as $G_{\eta,k} = VV'\, G_{\eta,k} + (I - VV')\, G_{\eta,k}$.

(ii) The influence function annihilates the bias-neutral component: $\Lambda\, G_{\eta,k} = U\Sigma\, w_k$, where $w_k = V'\, G_{\eta,k} \in \mathbb{R}^2$.

(iii) The squared bias sensitivity satisfies $\|\Lambda\, G_{\eta,k}\|^2 = \sigma_1^2\, w_{k,1}^2 + \sigma_2^2\, w_{k,2}^2$.

Proof. (i) Since V has orthonormal columns, $V'V = I_2$, so $P = VV'$ is an orthogonal projector ($P^2 = P$, $P' = P$). Then $G_{\eta,k} = P\, G_{\eta,k} + (I - P)\, G_{\eta,k}$, and the two components are orthogonal: $(P\, G_{\eta,k})'(I - P)\, G_{\eta,k} = G_{\eta,k}'\, P(I - P)\, G_{\eta,k} = 0$ since $P(I - P) = P - P^2 = 0$.

(ii) $\Lambda\, G_{\eta,k} = U\Sigma V'\, G_{\eta,k} = U\Sigma\, w_k$. In particular, $\Lambda\, (I - VV')\, G_{\eta,k} = U\Sigma V'(I - VV')\, G_{\eta,k} = U\Sigma(V' - V'VV')\, G_{\eta,k} = U\Sigma(V' - V')\, G_{\eta,k} = 0$, confirming that the bias-neutral component is annihilated.

(iii) Since U is orthogonal, $\|\Lambda\, G_{\eta,k}\|^2 = \|U\Sigma\, w_k\|^2 = \|\Sigma\, w_k\|^2 = \sigma_1^2\, w_{k,1}^2 + \sigma_2^2\, w_{k,2}^2$. □

Remark 4. The rank-two constraint and SVD decomposition are model-free: they depend only on the number of preference parameters (two) and the structure of the GMM influence function, not on the specific equilibrium model. The model enters exclusively through $G_{\eta,k}$, which determines how each structural parameter projects onto the bias-relevant subspace spanned by V.

Appendix F. Simulation Robustness

This appendix collects additional simulation diagnostics referenced in Section III.

A. *Profile Estimation*

The profile GMM estimator concentrates out λ for each trial (γ, ψ):

$$(\hat{\gamma}, \hat{\psi}) = \arg\min_{\gamma,\psi} \; Q\big(\gamma, \psi, \hat{\lambda}(\gamma, \psi)\big), \qquad \hat{\lambda}(\gamma, \psi) = \arg\min_{\lambda} \; Q(\gamma, \psi, \lambda). \tag{F1}$$

Table XI compares this profile approach with the baseline joint estimator. The profile estimator halves the cross-path dispersion of all three parameters while leaving the median estimates unchanged.

	$\hat{\gamma}$			$\hat{\psi}$			$\hat{\lambda}$		
Estimator	Median	Bias	Std	Median	Bias	Std	Median	Bias	Std
Oracle (true R_w)	11.4	+1.4	0.8	1.64	+0.14	0.06	–	–	–
Joint (baseline)	7.2	−2.8	1.4	1.42	−0.08	0.09	0.32	+0.02	0.016
Profile GMM	7.2	−2.8	**1.0**	1.42	−0.08	**0.07**	0.32	+0.02	**0.009**

Table XI: **Profile vs. Joint Estimation.** 50 paths × 2,000 years, monthly data. True: $\gamma = 10$, $\psi = 1.5$, $\lambda = 0.30$. Bold entries highlight dispersion reductions.

B. *Finite-Sample Diagnostics*

Table XII examines the behavior of the oracle and joint estimators as the sample size increases from 500 to 5,000 years. The oracle estimator's standard deviation shrinks from 5.7 to 0.5 for $\hat{\gamma}$ and from 0.23 to 0.04 for $\hat{\psi}$, confirming $\sqrt{T}$-consistency. The joint estimator displays a stable median bias of approximately -2.5 units for γ across all sample sizes, with precision improving at the parametric rate. The persistence of the median bias as T grows confirms that this residual discrepancy is a population-level phenomenon driven by the BKY log-linear approximation, not finite-sample noise.

C. *Robustness of Moment Conditions*

Table XIII varies the moment condition configuration to assess sensitivity. Dropping the wealth Euler equation leaves all estimates unchanged, indicating that the cross-sectional pricing restrictions alone suffice to identify λ when the dividend claims are sufficiently het-

	$\hat{\gamma}$			$\hat{\psi}$			$\hat{\lambda}$		
T (years)	Median	Bias	Std	Median	Bias	Std	Median	Bias	Std
Oracle (uses true R_w)									
500	11.3	+1.3	5.7	1.63	+0.13	0.23	–	–	–
2,000	11.4	+1.4	0.7	1.66	+0.16	0.06	–	–	–
5,000	11.2	+1.2	0.5	1.64	+0.14	0.04	–	–	–
Joint estimator									
500	7.8	−2.2	3.0	1.45	−0.05	0.22	0.32	+0.02	0.10
2,000	7.3	−2.7	2.4	1.44	−0.06	0.16	0.32	+0.02	0.05
5,000	7.6	−2.4	1.5	1.44	−0.06	0.10	0.32	+0.02	0.02

Table XII: **Finite-Sample Diagnostics.** 50 Monte Carlo paths per sample size. True: $\gamma = 10$, $\psi = 1.5$, $\lambda = 0.30$.

erogeneous. Expanding the cross-section from 5 to 25 BKY dividend claims tightens all estimates, with $\hat{\gamma}$ dispersion falling from 2.0 to 1.0. The J-test rejects at 100 percent for all joint configurations, reflecting the log-linear approximation error in the BKY model rather than estimator misspecification.

	$\hat{\gamma}$		$\hat{\psi}$		$\hat{\lambda}$		
Configuration	Median	Std	Median	Std	Median	Std	J-rej
Oracle (true R_w)	11.3	0.8	1.65	0.07	–	–	27%
Joint, 5 assets (baseline)	7.5	2.0	1.44	0.13	0.32	0.03	100%
Joint, 5 assets, no wealth Euler	7.5	2.0	1.44	0.13	0.32	0.03	100%
Joint, 25 assets	6.8	1.0	1.40	0.06	0.33	0.02	100%

Table XIII: **Robustness of the Joint Estimator.** T = 2,000 years monthly, 50 paths. True: $\gamma = 10$, $\psi = 1.5$, $\lambda = 0.30$.

D. Wald Test Power Across λ

Table XIV reports the power of the Wald test of $H_0 : \lambda = 1$ across structural wealth shares, using 62 years of monthly data. Power rises from 84 percent at $\lambda = 0.12$ to 98 percent at $\lambda = 0.50$, reflecting the stronger signal available when financial wealth constitutes a larger share of total wealth. The preference parameter estimates ($\hat{\gamma}_{\text{med}} \in [6.3, 7.0]$) are essentially invariant to the true λ, confirming that identification of preferences does not depend on the

wealth share.

	$\hat{\lambda}$		Wald Test		$\hat{\gamma}$	
λ_0	Median	Std	Power (5%)	SE($\hat{\lambda}$)	Median	Bias
0.12	0.154	0.292	84%	0.077	6.3	−3.7
0.20	0.231	0.250	87%	0.069	6.4	−3.6
0.30	0.326	0.211	87%	0.060	6.6	−3.4
0.50	0.517	0.041	98%	0.039	7.0	−3.0

Table XIV: **Wald Test Power Across Structural Wealth Shares.** 100 paths, 62 years monthly. True: $\gamma = 10$, $\psi = 1.5$.

E. Sensitivity to Model Calibration

Table XV varies the preference parameters and dividend leverage across seven calibrations. The market-proxy bias ($\hat{\gamma}^* \approx 66$–$73$) is remarkably stable across all specifications with $\gamma \geq 10$, confirming that the bias mechanism is universal. The joint estimator achieves 94–96 percent improvement in the baseline region but is less effective at extreme risk aversion ($\gamma = 5$ or $\gamma = 20$) or near-unit EIS ($\psi = 1.01$), where the optimization landscape flattens.

				Market Proxy		Joint Estimator			
Calibration	γ	ψ	θ	$\hat{\gamma}^*$	$\hat{\psi}^*$	$\hat{\gamma}$	$\hat{\psi}$	$\hat{\lambda}$	Improve
Baseline	10	1.50	−27	73.4	0.04	7.6	1.44	0.32	96%
Low RRA	5	1.50	−12	1.0	0.76	2.4	1.33	0.38	35%
High RRA	20	1.50	−57	66.3	0.00	0.0	1.02	0.11	57%
Near-unit EIS	10	1.01	−909	72.3	0.05	0.0	0.64	0.99	84%
High EIS	10	2.50	−15	70.0	0.04	6.4	2.00	0.35	94%
Low leverage	10	1.50	−27	71.8	0.04	7.7	1.45	0.32	96%
High leverage	10	1.50	−27	70.7	0.05	7.5	1.44	0.32	96%

Table XV: **Sensitivity to Model Calibration.** 50 paths, $T = 2{,}000$ years monthly, $\lambda = 0.30$.

F. Validation Under Exact Euler Equations

To definitively separate estimator bias from approximation error, we construct a discrete-state Markov chain economy where the equilibrium is solved exactly. The price-consumption

and price-dividend ratios are solved numerically to machine precision (10^{-12}).

	$\hat{\gamma}$		$\hat{\psi}$		$\hat{\lambda}$	
	Median	Bias	Median	Bias	Median	Bias
True values	10.0	–	1.50	–	0.30	–
BKY model (log-linear)						
Market proxy	73.7	+63.7	0.04	−1.46	–	–
Oracle (true R_w)	11.3	+1.3	1.65	+0.15	–	–
Joint estimator	7.5	−2.5	1.44	−0.06	0.32	+0.020
Exact Markov chain						
Joint (oracle start)	9.3	−0.7	1.52	+0.02	0.30	+0.003

Table XVI: **Exact Model Validation.** BKY: 50 paths, $T = 2{,}000$ years monthly. Exact model: 30 paths, $T = 2{,}000$ years. True: $\gamma = 10$, $\psi = 1.5$, $\lambda = 0.30$.

The joint estimator recovers the structural parameters with negligible bias under the exact model: $\hat{\gamma}_{\text{med}} = 9.3$, $\hat{\psi}_{\text{med}} = 1.52$, $\hat{\lambda}_{\text{med}} = 0.30$. This confirms that the persistent $\sim$2.5-unit γ bias in the BKY simulations is entirely attributable to the log-linear approximation error, not to any deficiency of the estimator.

G. Accuracy of the Wealth Decomposition

The joint estimator and the bias correction formulas both rely on the portfolio decomposition $R_w = \lambda R_m + (1-\lambda) R_h$ (Assumption 3). In gross returns, this identity is exact by construction: given the true wealth return R_w and any structural $\lambda \in (0, 1)$, the non-tradable return $R_h = (R_w - \lambda R_m)/(1-\lambda)$ satisfies the decomposition identically. The log-return counterpart $r_w = \lambda r_m + (1-\lambda) r_h$ is an approximation with a Jensen's inequality error of order $O(\sigma^2)$. We verify this numerically in the BKY economy using 100 independent paths of 2,000 years of monthly data for three structural wealth shares ($\lambda \in \{0.12, 0.30, 0.50\}$). The gross-return identity holds to machine precision ($\sim 10^{-16}$) on every path. At the benchmark calibration ($\lambda = 0.30$), the log-return approximation error has a median mean absolute value of 0.04 percent at monthly frequency, two orders of magnitude smaller than the bias effects documented in Section III. Even at $\lambda = 0.50$ the median error is only 0.10 percent. The R^2 of $\log(R_w)$ on $\log(R_m)$ alone is approximately 0.51 –the market return explains about half the variation in log wealth returns despite the large volatility gap ($\sigma(r_w)/\sigma(r_m) = 0.17$). The

backed-out human capital return has negative correlation with market returns ($\hat{\rho} \approx -0.84$ at $\lambda = 0.30$), reflecting the residual construction: since the wealth return is much less volatile than the market, the non-tradable component must offset market fluctuations.

These results confirm that the portfolio decomposition introduces negligible approximation error at the frequencies relevant for estimation, and that the theoretical distinction between log and gross λ (Remark 1) is empirically inconsequential in the BKY economy.

Appendix G. Empirical Estimation Details

This appendix provides the empirical estimation details underlying Figure 1. The figure plots nonlinear GMM estimates of the coefficient of relative risk aversion (γ) and the elasticity of intertemporal substitution (ψ) across multiple preference specifications and consumption measures.

A. Data and Consumption Measures

Real personal consumption expenditure data come from the U.S. National Income and Product Accounts (NIPA), Bureau of Economic Analysis. Municipal waste generation data are from the Environmental Protection Agency. Equity portfolio returns and factors are from Kenneth French's website. The sample spans 1961–2022.

Five alternative measures of consumption growth are considered:

1. **NIPA.** One-year growth in annual real non-durables and services per capita. This is the standard measure underlying the equity premium puzzle; its low volatility ($\sigma = 1.59\%$) is central to Mehra and Prescott's (1985) argument.

2. **Ultimate.** Three-year cumulative consumption growth from Parker and Julliard (2005), designed to capture slow adjustment of consumption to asset return shocks.

3. **Q4-Q4.** Year-on-year growth in fourth-quarter consumption from Jagannathan and Wang (2007). This measure exploits the observation that fourth-quarter consumption growth is more strongly correlated with annual returns than calendar-year growth, possibly due to holiday-related discretionary spending.

4. **Unfiltered.** Real per capita consumption growth constructed following Kroencke (2017), which avoids the smoothing embedded in standard NIPA revisions. Unfiltered consumption is substantially more volatile than NIPA ($\sigma \approx 2.5\%$), potentially better capturing the consumption risk that equities compensate.

5. **Garbage.** One-year growth in total municipal solid waste generation less yard trimmings from Savov (2011). Garbage proxies for the flow of consumption services from

all goods, including durables, and is measured independently of the national accounts. It has the highest correlation with the market return ($\rho \approx 0.57$) among all measures.

All measures are adjusted for U.S. population growth. The test assets are the risk-free rate, excess returns on the aggregate stock market, and excess returns on the 25 Fama and French (1992) size and book-to-market portfolios, deflated using the non-durables and services price index. Consumption is aligned with returns using a beginning-of-period timing restriction as in Campbell (2003), matching returns during time t with consumption growth from t to $t+1$.

B. Estimation Setup

Preference specifications

Three preference specifications are estimated by nonlinear GMM, ordered by increasing generality.

CC1 (Power utility). The stochastic discount factor under constant relative risk aversion is

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}. \tag{G1}$$

The single free parameter is γ; the EIS is $\psi = 1/\gamma$ by construction. CC1 is the canonical Lucas (1978)–Breeden (1979) specification underlying the original equity premium puzzle of Mehra and Prescott (1985). Because (G1) does not depend on $R_{w,t+1}$, CC1 serves as a natural benchmark: it is immune to the wealth proxy problem that drives the bias in recursive-utility models. All CC1 estimates lie on the $\psi = 1/\gamma$ curve in Figure 1.

CC2 (Recursive utility). Under the Epstein and Zin (1989)–Weil (1990) specification, the SDF is

$$M_{t+1} = \beta^{\theta} \left(\frac{C_{t+1}}{C_t} \right)^{-\theta/\psi} R_{w,t+1}^{\theta-1}, \tag{G2}$$

where $\theta = (1-\gamma)/(1-1/\psi)$ and $R_{w,t+1}$ is the gross return on total wealth, proxied by the market return $R_{m,t+1}$. Both γ and ψ are estimated as free parameters. CC2 nests CC1 as the special case $\gamma = 1/\psi$ ($\theta = 1$), where $R_{w,t+1}$ drops out of the SDF.

CC3 (Durable goods utility). The Yogo (2006) model extends recursive preferences

to non-separable utility over non-durable consumption C_t and durable consumption D_t. The intratemporal aggregator is CES:

$$V_t = \left[(1-\alpha) + \alpha \left(\frac{C_t}{D_t} \right)^{1-1/\eta} \right]^{1/(1-1/\eta)},$$

where α is the non-durable share and η is the intratemporal elasticity of substitution between goods. The SDF becomes

$$M_{t+1} = \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-1/\psi} \left(\frac{V_{t+1}}{V_t} \right)^{1/\eta - 1/\psi} R_{w,t+1}^{1-1/\theta} \right]^{\theta},$$

with γ and ψ estimated freely while $\alpha = 0.827$, $\eta = 0.790$, and the durable depreciation rate $\delta = 0.21$ are calibrated following Yogo (2006). CC3 introduces an additional moment condition from the durable goods Euler equation, which requires the relative price of durables to non-durables.

Moment conditions

For each specification, estimation exploits three sets of moment conditions with increasing identifying content.

Time-series moments (TS). The risk-free rate Euler equation $\mathbb{E}[M_{t+1} R_{f,t}] = 1$ and the market excess return pricing restriction $\mathbb{E}[M_{t+1} R^e_{m,t+1}] = 0$ provide two moment conditions. Under CC2, a third restriction comes from the wealth Euler equation (equation (6) in the main text):

$$\mathbb{E}\left[\frac{1}{\theta} \left(\left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-1/\psi} R_{w,t+1} \right)^{\theta} - 1 \right) \right] = 0. \tag{G3}$$

This restriction disciplines the joint dynamics of consumption growth and the wealth return without involving any additional test asset. The TS configuration thus provides 2 moments for CC1 and 3 moments for CC2.

Cross-sectional moments (CS). The pricing errors on $N = 25$ Fama and French (1992)

size and book-to-market portfolios:

$$\mathbb{E}[M_{t+1}R^e_{j,t+1}] = 0, \quad j = 1, \ldots, 25,$$

supply 25 moment conditions. This configuration tests the model's ability to explain the cross-section of expected returns.

Joint moments (TS+CS). The union of TS and CS moment conditions. For CC1, this yields $2 + 25 = 27$ moments with 1 parameter (γ); for CC2, $3 + 25 = 28$ moments with 2 parameters (γ, ψ); and for CC3, $3+25+1 = 29$ moments with 2 parameters (γ, ψ), where the additional moment is the durable goods Euler equation. Figure 1 reports the joint TS+CS estimates, which provide the most demanding test of each specification.

GMM implementation

All models are estimated by one-step GMM with the identity matrix as the weighting matrix ($W = I$), following the recommendation of Hansen and Singleton (1982) for potentially misspecified models where efficient two-step weighting may amplify finite-sample bias. The long-run covariance matrix is estimated using the Newey and West (1987) HAC estimator with 5 Newey–West lags (annual data). The discount factor is fixed at $\beta = 0.95$ throughout.

Parameter bounds are $\gamma \in [0.001, 5000]$ and $\psi \in [0.0001, 3]$. For CC1 and CC2, optimization uses constrained nonlinear minimization. For CC3, we use global optimization due to the additional nonlinearity introduced by the CES aggregator and the durable goods Euler equation. Starting values are chosen on a specification-by-specification basis: $\gamma_0 = 100$ for CC1; $(\gamma_0, \psi_0) = (100, 0.5)$ for CC2; $(\gamma_0, \psi_0) = (70, 0.5)$ for CC3, with consumption-measure-specific adjustments where convergence requires it.

Diagnostics. For each specification, we report: (i) point estimates and t-statistics based on HAC standard errors; (ii) the Hansen–Jagannathan distance (Hansen and Jagannathan, 1997) as a measure of model fit; and (iii) the HJ distance p-value. The J-statistic for the overidentifying restrictions test is also computed but not reported in the interest of space; all specifications are rejected at conventional levels.

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	66.40	77.96	104.37	17.33	23.07
(t-stat)	(2.12)	(3.21)	(2.72)	(1.10)	(2.95)
HJ dist	3.96	1.48	4.05	2.85	1.24
(p-value)	(0.00)	(0.24)	(0.00)	(0.01)	(0.03)

(a) CC1

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	91.77	78.29	130.76	26.21	34.56
(t-stat)	(3.94)	(3.29)	(4.25)	(3.01)	(2.53)
ψ	0.00	0.03	0.00	0.00	0.07
(t-stat)	(0.01)	(1.15)	(0.01)	(0.00)	(1.92)
HJ dist	3.93	1.45	3.64	1.89	1.21
(p-value)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)

(b) CC2

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	84.60	48.19	121.56	6.30	30.17
(t-stat)	(3.24)	(0.79)	(2.60)	(0.38)	(0.84)
ψ	0.01	0.05	0.01	0.22	0.10
(t-stat)	(0.52)	(1.94)	(0.92)	(0.62)	(3.74)
HJ dist	5.86	6.15	6.39	3.55	3.00
(p-value)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

(c) CC3

Table XVII: **Estimation of non-linear models**

Estimated preference parameters across consumption models and measures. Test assets consists of the market excess return, the risk-free rate, and the 25 Fama and French (1992) size and book-to-market portfolios. Panel (a) reports results for the CC1 model, panel (b) for the CC2, and panel (c) for the CC3 model. The sample spans 1961-2022, with Garbage ending in 2018.

C. *Results by Preference Specification*

Table XVII reports the joint TS+CS estimates. Under power utility (CC1), NIPA consumption yields $\hat{\gamma} = 66$ (t-stat $= 2.12$), and alternative measures yield $\hat{\gamma}$ ranging from 17 (Unfiltered) to 104 (Q4-Q4), with implied EIS between 0.01 and 0.06. No measure produces economically plausible risk aversion. Under recursive utility (CC2), the pattern persists: NIPA yields $\hat{\gamma} = 92$, Q4-Q4 yields $\hat{\gamma} = 131$, while Unfiltered ($\hat{\gamma} = 26$) and Garbage ($\hat{\gamma} = 35$) produce more moderate but still elevated estimates. The EIS is uniformly near zero ($\hat{\psi} \in [0.00, 0.20]$), with only Garbage showing marginal evidence of a positive EIS ($\hat{\psi} = 0.07$, t-stat $= 1.92$). Under durable goods utility (CC3), greater flexibility does not resolve the puzzle: Unfiltered consumption produces the most reasonable estimates ($\hat{\gamma} = 6.3$, $\hat{\psi} = 0.22$), but other measures yield $\hat{\gamma}$ from 30 (Garbage) to 122 (Q4-Q4), and all specifications are rejected by the J-test.

D. *Robustness Across Moment Configurations*

As a robustness check, each preference specification is estimated under the time-series (TS), cross-sectional (CS), and joint (TS+CS) moment configurations separately. Table XVIII reports the time-series estimates and Table XIX the cross-sectional estimates. Under CC1, risk aversion estimates range from $\hat{\gamma} = 17$ (Unfiltered, CS) to $\hat{\gamma} = 104$ (Q4-Q4, TS+CS), all well above plausible values. Under CC2, the time-series moments alone yield $\hat{\gamma}$ from 26 (Unfiltered) to 131 (Q4-Q4), while the EIS is uniformly near zero ($\hat{\psi} \in [0.00, 0.20]$); cross-sectional moments produce similar distortions, confirming that the near-zero EIS is a robust feature of recursive-utility estimation rather than an artifact of a particular moment condition. Under CC3, Unfiltered consumption yields the most plausible estimates ($\hat{\gamma} = 6.3$, $\hat{\psi} = 0.22$), but other measures produce $\hat{\gamma}$ from 30 (Garbage) to 122 (Q4-Q4), with HJ distances generally larger than for CC1 and CC2, reflecting the burden of the additional durable goods moment. The invariance of the pathological pattern across moment configurations, preference specifications, and consumption measures is the central empirical motivation for the theoretical analysis in the main text.

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	90.78	79.37	130.00	26.14	26.92
(t-stat)	(3.77)	(3.26)	(4.01)	(2.86)	(3.29)
HJ dist	1.32	0.07	1.29	0.68	0.36
(p-value)	(0.00)	(0.72)	(0.01)	(0.01)	(0.11)

(a) CC1

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	91.78	79.13	130.76	26.20	40.20
(t-stat)	(3.89)	(3.28)	(4.22)	(2.82)	(3.25)
ψ	0.00	0.03	0.00	0.00	0.09
(t-stat)	(0.01)	(1.34)	(0.01)	(0.27)	(3.05)
HJ dist	1.19	0.22	1.10	0.42	0.00
(p-value)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

(b) CC2

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	84.54	48.14	121.63	21.98	30.03
(t-stat)	(0.01)	(0.03)	(0.01)	(0.02)	(0.05)
ψ	0.01	0.05	0.01	0.02	0.10
(t-stat)	(0.00)	(0.31)	(0.01)	(0.01)	(1.91)
HJ dist	5.82	6.14	6.36	2.59	2.37
(p-value)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

(c) CC3

Table XVIII: **Estimation of non-linear models: Time-series tests**
Estimated preference parameters across consumption models and measures. Test assets consists of the market excess return and the risk-free rate. Panel (a) reports results for the CC1 model, panel (b) for the CC2, and panel (c) for the CC3 model. The sample spans 1961-2022, with Garbage ending in 2018.

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	43.64	39.67	59.90	15.65	22.12
(t-stat)	(2.58)	(1.99)	(3.24)	(1.76)	(2.82)
HJ dist	1.60	0.36	1.31	2.50	1.10
(p-value)	(0.10)	(0.17)	(0.18)	(0.02)	(0.00)

(a) CC1

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	71.04	68.18	51.58	26.20	30.14
(t-stat)	(3.92)	(3.06)	(3.59)	(2.97)	(2.98)
ψ	0.20	0.06	0.04	0.00	0.07
(t-stat)	(5.52)	(1.40)	(4.59)	(0.01)	(1.92)
HJ dist	153.03	1.27	2.13	1.85	1.14
(p-value)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

(b) CC2

	NIPA	Ultimate	Q4-Q4	Unfiltered	Garbage
γ	84.60	48.12	121.62	22.16	30.14
(t-stat)	(2.90)	(1.92)	(2.08)	(2.74)	(1.08)
ψ	0.01	0.05	0.01	0.02	0.10
(t-stat)	(0.53)	(2.68)	(0.92)	(0.35)	(3.75)
HJ dist	0.66	0.33	0.66	1.58	1.86
(p-value)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

(c) CC3

Table XIX: **Estimation of non-linear models: Cross-sectional tests**
Estimated preference parameters across consumption models and measures. Test assets consists of the 25 Fama and French (1992) size and book-to-market portfolios. Panel (a) reports results for the CC1 model, panel (b) for the CC2, and panel (c) for the CC3 model. The sample spans 1961-2022, with Garbage ending in 2018.

E. The Pattern in Figure 1

Figure 1 plots all 15 estimation points in the (ψ, γ) plane. Circles denote power utility (CC1, constrained to the $\psi = 1/\gamma$ curve), squares denote recursive utility (CC2), and triangles denote durable goods utility (CC3). The calibrated values from Bansal and Yaron (2004) ($\gamma = 7.5$, $\psi = 1.5$) and Bansal, Kiku, and Yaron (2007) ($\gamma = 10$, $\psi = 1.5$) are shown for comparison.

Two features are immediately apparent. First, all estimates cluster in the upper-left region of the parameter space: high risk aversion ($\hat{\gamma} \gg 10$) and near-zero EIS ($\hat{\psi} \approx 0$). No combination of preference specification and consumption measure moves the estimates into the economically reasonable region near the calibrated values. Second, the clustering is remarkably tight given the heterogeneity in consumption measures –from NIPA's smooth national accounts data to Garbage's independently measured waste flows –suggesting a common underlying cause rather than model-specific misspecification.

F. Connection to Theory

The theoretical framework of Section II predicts exactly this pattern. When the return on wealth is proxied by the market return, the full Euler pseudo-true parameters (Proposition 4) map any true (γ, ψ) in the plausible region ($\gamma \in [5, 20]$, $\psi \in [0.5, 2]$) onto a narrow pathological band ($\gamma^* \in [98, 111]$, $\psi^* \approx 0$). With empirically relevant wealth shares ($\lambda \approx 0.1$–0.3), true risk aversion of $\gamma = 10$ produces $\gamma^* \approx 70$–100 –precisely where the data places these estimates.

The near-invariance result (Section D) implies that the observed estimates are essentially uninformative about the true risk aversion: the mapping from true to pseudo-true parameters compresses all plausible inputs into a narrow band of outputs. That this compression is universal across consumption measures is consistent with the mechanism: wealth mismeasurement biases the return component of the SDF, and no improvement in consumption measurement can address a return measurement problem.