The Belief-Demand Dichotomy in Cross-Sectional Asset Pricing*

[Author Name]†

March 20, 2026

Abstract

Characteristic-sorted equity portfolios carry unpriced variance components ranging from 30% to 99% of total variance. What determines this share? A continuous-time economy with Epstein-Zin preferences, Markov regime switching, institutional mandates, and stochastic volatility delivers a sharp dichotomy. Investor beliefs about systematic factors have a quantitatively negligible effect on the unpriced variance of characteristic-sorted portfolios: the invariance holds exactly when investors share a common risk aversion and approximately (deviation $< 1\%$) when risk aversions differ. Three channels break the invariance. Institutional demand mandates inject non-fundamental variance proportional to mandate-shock volatility. Stochastic characteristic-specific volatility creates price variation through the discount-rate channel, with a surprising stabilization result: disagreement about volatility *reduces* this component. Wealth redistribution from heterogeneous risk aversion contributes at second order. The relative importance of these channels varies across characteristics: mandate-driven variance dominates for characteristics with large institutional clienteles, while the stochastic-volatility channel dominates for characteristics with low institutional involvement.

JEL codes: G11, G12, G23.

Keywords: factor models, unpriced risk, belief heterogeneity, institutional mandates, stochastic volatility, characteristic portfolios.

*The author thanks [acknowledgments to be added]. All errors are the author's own.

†[Affiliation and contact information.]

1 Introduction

Why do characteristic-sorted equity portfolios carry so much risk that no factor model prices? Chernov et al. [2025] document that 30–99% of the variance of standard long-short factor portfolios is orthogonal to the stochastic discount factor. Hedging this unpriced component more than doubles Sharpe ratios. The magnitudes are large, vary across characteristics, and lack a theoretical explanation.

This paper provides an explanation. In a continuous-time economy with Epstein-Zin preferences, Markov regime switching, institutional demand mandates, and stochastic volatility, a dichotomy governs the unpriced variance share. Investor beliefs about systematic factors have a quantitatively negligible effect on the variance of characteristic-sorted portfolios. Institutional demand frictions and stochastic volatility are the economically relevant channels, and their relative importance varies across characteristics.

The model features two Lucas trees: a systematic asset M and a characteristic-specific asset C with independent dividend processes. Two investor types have Duffie-Epstein stochastic differential utility, potentially differing in risk aversion and elasticity of intertemporal substitution. A two-state Markov chain governs the business cycle. Three additional state variables capture the channels that break the invariance: institutional mandate targets following an Ornstein-Uhlenbeck process, stochastic volatility of C's dividends via a CIR process, and the endogenous wealth distribution.

Six nested cases progressively activate these features. Cases 1 and 2 establish the invariance result. With common risk aversion, the instantaneous return variance of the characteristic asset equals its fundamental dividend variance, regardless of the elasticity of intertemporal substitution, regime, or properties of the systematic asset. Per-dollar demand for C depends only on the ratio of risk premium to variance, and in equilibrium both scale identically with fundamental volatility. The portfolio weight on C is pinned by risk aversion alone. Cases 3 through 6 identify three breaking channels and assemble the full variance decomposition in Theorem 8.

Three main results emerge.

The invariance holds exactly under common risk aversion with arbitrary EIS heterogeneity and arbitrary regime switching (Propositions 1 and 2). When risk aversions differ, the deviation is second-order in a bounded heterogeneity parameter and quantitatively below 1% for empirically plausible values (Proposition 3). The invariance extends to smooth ambiguity preferences, private signals in a noisy rational expectations equilibrium, and bilateral search frictions (Appendix B).

Two first-order channels break the invariance. Institutional mandates inject non-fundamental variance V_M proportional to mandate-shock volatility, through a demand-pressure channel orthogonal to both systematic and fundamental risk (Proposition 4). Countercyclical S_U requires mandate-shock amplification to exceed the ratio of systematic volatilities across regimes (Proposition 5). Stochastic characteristic-specific volatility creates price variation V_v through the discount-rate response of the price-dividend ratio (Proposition 6). A stabilization result accompanies the volatility channel: when investors disagree about characteristic-specific volatility, the market price responds less than one-for-one, reducing this variance component (Proposition 7).

The channel ordering is characteristic-dependent. For characteristics with large institutional clienteles, mandate-driven variance dominates. For characteristics with low institutional involvement, the stochastic-volatility channel dominates. Wealth redistribution from heterogeneous risk aversion is always the smallest channel by at least an order of magnitude.

The paper contributes to several literatures. The Chernov et al. [2025] finding calls for a structural explanation of why characteristic factors embed large unpriced components. This paper traces the unpriced share to institutional demand primitives and stochastic volatility rather than to beliefs or information frictions. The demand-system approach of Koijen and Yogo [2019] emphasizes heterogeneous institutional demand as a price driver; the present model formalizes one channel through which demand heterogeneity generates non-fundamental variance. The preferred-habitat framework of Vayanos and Vila [2021] shows how clientele demand and limited arbitrage generate non-fundamental price variation in bond markets; this paper adapts the logic to equity characteristics, adds stochastic volatility as a competing channel, and derives cross-characteristic predictions absent from the bond-market framework. The recursive-utility literature following Duffie and Epstein [1992] and Bansal and Yaron [2004] emphasizes the role of the EIS in portfolio choice; the invariance result shows that this parameter is irrelevant for the variance composition of characteristic portfolios when investors share a common risk aversion. Basak and Cuoco [1998] study wealth redistribution under heterogeneous risk aversion; Proposition 3 shows that this channel contributes at second order to characteristic-specific variance, quantitatively negligible relative to the mandate and stochastic-volatility channels.

Section 2 presents the model. Section 3 establishes the invariance result. Section 4 identifies the breaking channels and derives the full variance decomposition. Section 5 calibrates the model. Section 6 presents testable predictions. Section 7 concludes.

2 Model

2.1 Assets

The economy has three investment opportunities in continuous time.

A *risk-free asset* pays instantaneous rate $r(s_t)$, potentially dependent on the business-cycle regime s_t, determined in equilibrium.

Asset M (systematic) is a Lucas tree in unit supply with dividend process

$$\frac{\mathrm{d}D_M}{D_M} = \mu_D^M(s_t)\,\mathrm{d}t + \sigma_D^M(s_t)\,\mathrm{d}W_F, \tag{1}$$

where W_F is a standard Brownian motion. Both drift and diffusion depend on the regime s_t.

Asset C (characteristic-specific) is a Lucas tree in unit supply with dividend process

$$\frac{\mathrm{d}D_C}{D_C} = \mu_D^C\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_u, \tag{2}$$

where W_u is a standard Brownian motion independent of W_F, and v_t is the instantaneous dividend variance. In Cases 1–4, $v_t = \sigma_C^2$ is a constant. In Cases 5–6, v_t follows a CIR process (Section 4.3).

Remark 1. *The independence $\langle W_F, W_u\rangle_t = 0$ defines what "characteristic-specific" means: asset C's cash flows are orthogonal to all systematic factors, including the business cycle. The theory asks whether equilibrium forces (preferences, heterogeneity, institutional frictions) can create state-dependent return variance for an asset whose fundamental variance is state-independent.*

2.2 Business cycle

The state $s_t \in \{G, B\}$ follows a continuous-time Markov chain with transition intensities λ_{GB} (from expansion to recession) and λ_{BG} (from recession to expansion). The regime-switch process is independent of W_F and W_u.

2.3 Investors

Two types $j \in \{H, L\}$ populate the economy, with initial wealth $W_0^j > 0$. Total wealth is $W_t = W_t^H + W_t^L$. The wealth share is $\omega_t = W_t^H / W_t$.

Each type has stochastic differential utility [Duffie and Epstein, 1992]:

$$V_t^j = \mathbb{E}_t\left[\int_t^\infty f_j(c_\tau^j, V_\tau^j)\,\mathrm{d}\tau\right], \tag{3}$$

with the normalized CES-Kreps-Porteus aggregator

$$f_j(c,v) = \frac{\beta_j}{1-1/\psi_j}\cdot\frac{c^{1-1/\psi_j} - \big((1-\gamma_j)v\big)^{1/\theta_j}}{\big((1-\gamma_j)v\big)^{1/\theta_j-1}}, \tag{4}$$

where $\theta_j = (1-\gamma_j)/(1-1/\psi_j)$, $\gamma_j > 0$ is the coefficient of relative risk aversion, $\psi_j > 0$ is the elasticity of intertemporal substitution, and $\beta_j > 0$ is the rate of time preference. When $\gamma_j = 1/\psi_j$, the aggregator reduces to time-additive CRRA.

Type j chooses consumption rate $c_t^j \geq 0$ and portfolio weights (π_M^j, π_C^j) on assets M and C to maximize V_0^j subject to the budget constraint

$$\frac{\mathrm{d}W^j}{W^j} = \left[r + \pi_M^j(\mu_M - r) + \pi_C^j(\mu_C - r) - c^j/W^j\right]\mathrm{d}t + \pi_M^j\sigma_M\,\mathrm{d}W_F + \pi_C^j\sigma_C^{eq}\,\mathrm{d}W_u, \tag{5}$$

where μ_M, μ_C are equilibrium expected returns and σ_M, σ_C^{eq} are equilibrium return volatilities.

2.4 Mandates (Cases 4 and 6)

Type L faces a soft mandate constraint on characteristic exposure. The mandate target τ_t follows

$$\mathrm{d}\tau_t = \kappa(\bar{\tau}(s_t) - \tau_t)\,\mathrm{d}t + \sigma_\tau(s_t)\,\mathrm{d}W_\tau, \tag{6}$$

where W_τ is independent of W_F, W_u, and the regime process, $\kappa > 0$ governs mean reversion, $\bar{\tau}(s_t)$ is the regime-dependent long-run target, and $\sigma_\tau(s_t)$ is the regime-dependent mandate-shock volatility. The mandate enters Type L's Bellman equation as a quadratic penalty:

$$-\frac{\chi}{2}\frac{(W^L)^{1-\gamma_L}}{1-\gamma_L}(\pi_C^L - \tau_t)^2, \tag{7}$$

where $\chi > 0$ measures mandate strictness. The multiplicative scaling preserves homotheticity.

2.5 Stochastic volatility (Cases 5 and 6)

The instantaneous variance v_t of C's dividends follows a CIR process:

$$\mathrm{d}v_t = \kappa_v(\bar{v} - v_t)\,\mathrm{d}t + \sigma_v\sqrt{v_t}\,\mathrm{d}W_v, \tag{8}$$

where W_v is independent of W_F, W_u, and W_τ, $\bar{v}$ is the long-run mean variance, $\kappa_v > 0$ is the speed of mean reversion, and $\sigma_v > 0$ is the volatility of volatility. The Feller condition $2\kappa_v\bar{v} > \sigma_v^2$ ensures $v_t > 0$.

2.6 Equilibrium

Definition 1. *A* competitive equilibrium *consists of price processes* (P_M, P_C, r)*, consumption plans* (c^H, c^L)*, and portfolio policies* $(\pi_M^j, \pi_C^j)_{j\in\{H,L\}}$ *such that: (i) each type maximizes* (3) *subject to* (5) *(and the mandate penalty* (7) *for Type L when active); (ii) goods markets clear,* $c^H + c^L = D_M + D_C$*; and (iii) asset markets clear,* $\omega\pi_i^H + (1-\omega)\pi_i^L = w_i$ *for* $i \in \{M, C\}$*, where* $w_i = P_i/(P_M + P_C)$ *is the capitalization weight.*

2.7 State space and nesting

The full model has state vector $(s_t, \omega_t, v_t, \tau_t)$. Table 1 summarizes which features are active in each case.

Table 1: Six nested cases

Feature	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6
Risk aversion γ	Common	Common	Hetero.	Hetero.	Common	Hetero.
EIS ψ	Common	Hetero.	Hetero.	Hetero.	Common	Hetero.
Regime switching	No	Yes	Yes	Yes	Yes	Yes
Mandates (τ_t)	No	No	No	Yes	No	Yes
Stochastic vol (v_t)	No	No	No	No	Yes	Yes

3 The Invariance Result

3.1 Case 1: Common preferences, constant opportunities

Both investor types share (γ, ψ, β) and no regime switching operates ($\lambda_{GB} = \lambda_{BG} = 0$). Opportunities are constant, and a representative agent holds all wealth.

Conjecture the value function $V = AW^{1-\gamma}/(1-\gamma)$ with constant $A > 0$. The Bellman equation for Duffie-Epstein utility is

$$0 = \sup_{c, \pi_M, \pi_C} f(c, V) + \mathcal{A}V, \tag{9}$$

where $\mathcal{A}V$ is the infinitesimal generator under the controlled wealth process. Substituting and differentiating with respect to portfolio weights yields the first-order conditions:

$$\pi_M^* = \frac{\mu_M - r}{\gamma \sigma_M^2}, \qquad \pi_C^* = \frac{\mu_C - r}{\gamma \sigma_C^2}. \tag{10}$$

The EIS ψ does not appear. With constant opportunities, hedging demand is zero and the portfolio rule is purely myopic, consistent with Duffie and Epstein [1992] and Skiadas [2009].

Market clearing $\pi_C^* = w_C$ gives the equilibrium risk premium $\mu_C - r = \gamma w_C \sigma_C^2$. Since all parameters are constants, the price-dividend ratio $Q_C = P_C / D_C$ is constant, and the return on C is

$$\mathrm{d}R_C = \left(\frac{1}{Q_C} + \mu_D^C \right) \mathrm{d}t + \sigma_C \, \mathrm{d}W_u. \tag{11}$$

Proposition 1 (Invariance: common preferences, no regime switching)**.** *In the Duffie-Epstein economy with common (γ, ψ, β) and $\lambda_{GB} = \lambda_{BG} = 0$:*

(a) The optimal portfolio weight on C is $\pi_C^* = (\mu_C - r)/(\gamma \sigma_C^2)$, independent of ψ.

(b) The equilibrium risk premium on C is $\mu_C - r = \gamma w_C \sigma_C^2$, a constant.

(c) The instantaneous return variance is $\mathrm{Var}(\mathrm{d}R_C)/\mathrm{d}t = \sigma_C^2$.

(d) The return variance of C depends only on the dividend volatility and is invariant to (γ, ψ, β), to the properties of asset M, and to the risk-free rate.

Proof. Parts (a)–(b) follow from the first-order conditions (10) and market clearing. Part (c): with constant Q_C, the return (11) has diffusion coefficient σ_C, so $\mathrm{Var}(\mathrm{d}R_C)/\mathrm{d}t = \sigma_C^2$. Part (d): $\sigma_C^2 = (\sigma_D^C)^2$ is a primitive of the dividend process. □

3.2 Case 2: Common γ, heterogeneous ψ, regime switching

Now $\lambda_{GB} > 0$, $\lambda_{BG} > 0$, and $\psi_H \neq \psi_L$. Investment opportunities are stochastic through the regime. Conjecture $V^j(W^j, s) = A^j(s)(W^j)^{1-\gamma}/(1-\gamma)$.

The Bellman equation acquires a jump term from regime transitions:

$$0 = \sup_{c^j, \pi_M^j, \pi_C^j} f_j(c^j, V^j) + \mathcal{A}^{\mathrm{diff}} V^j + \lambda_s \big[V^j(W^j, \bar{s}) - V^j(W^j, s) \big], \tag{12}$$

where $\bar{s}$ denotes the other regime and λ_s is the transition intensity from state s.

Portfolio first-order conditions. The jump term $\lambda_s[A^j(\bar{s}) - A^j(s)](W^j)^{1-\gamma}/(1-\gamma)$ does not depend on portfolio weights because wealth is continuous at regime transitions and the transition intensity is exogenous. The Brownian-driven portfolio conditions are therefore identical to Case 1:

$$\pi^j_M(s) = \frac{\mu_M(s) - r(s)}{\gamma \sigma_M(s)^2}, \qquad \pi^j_C(s) = \frac{\mu^{eq}_C(s) - r(s)}{\gamma (\sigma^{eq}_C)^2}. \tag{13}$$

No hedging demand appears. The Duffie-Epstein correction to the portfolio condition requires nonzero covariance between asset returns and the innovation in continuation utility. The regime-switch innovation is a compensated Poisson process, independent of W_F and W_u. The covariance is zero, so no hedging demand arises in the diffusion channel.

Identical portfolios. Since γ is common, both types hold the same portfolio per dollar of wealth. Market clearing for C gives $\mu^{eq}_C - r = \gamma w_C (\sigma^{eq}_C)^2$, where $w_C(s) = P_C/(P_M(s) + P_C)$.

Price of C. The risk premium on C under the Duffie-Epstein SDF satisfies $\mu_C - r = \gamma \sigma_{C,W} + (1-\theta)\sigma_{C,A}$, where $\sigma_{C,W}$ is the covariance of C's return with wealth and $\sigma_{C,A}$ is the covariance with the innovation in $\log A(s)$. The second term vanishes because C's diffusion return (driven by W_u) is independent of the Poisson regime-switch innovation in $A(s)$. The price-dividend ratio Q_C may depend on s through $w_C(s)$, but this dependence affects only the drift and the jump variance at regime switches, not the diffusion coefficient.

Proposition 2 (Invariance: common γ, heterogeneous ψ, regime switching)**.** *In the Duffie-Epstein economy with common γ, $\psi_H \neq \psi_L$, and two-state Markov regime switching:*

(a) Both types hold identical portfolio weights, equal to the myopic demands (13)*, for all $s \in \{G, B\}$.*

(b) Hedging demand for both assets is zero.

(c) The diffusion variance of C's return is $\mathrm{Var}^{\mathrm{diff}}(\mathrm{d}R_C)/\mathrm{d}t = \sigma_C^2$*, constant across regimes.*

(d) The return variance of C is invariant to ψ_H, ψ_L, the transition intensities λ_{GB} and λ_{BG}, the dividend parameters of M, and the wealth share ω_t.

Proof. Part (a): the first-order conditions (13) depend only on γ, which is common. Part (b): the regime innovation is a Poisson process independent of W_F and W_u, so $\mathrm{Cov}(\mathrm{d}R_i, \mathrm{d}\log A^j) =$

0. Part (c): with $\sigma_{C,A} = 0$, the price of C reflects only the myopic risk premium. Since C's dividends are state-independent and the risk premium on C (through the diffusion channel) does not depend on s, the diffusion coefficient of R_C is σ_C. Part (d): σ_C^2 is a primitive. $\square$

Why ψ-heterogeneity is irrelevant. The EIS governs how aggressively investors substitute consumption across time. Types H and L accumulate wealth at different rates, so ω_t drifts. But per-dollar demand for C depends only on γ. With common γ, aggregate demand for C per dollar of aggregate wealth is independent of how wealth is distributed. The wealth share ω cancels from the market-clearing condition.

4 What Breaks the Invariance

Three channels generate non-fundamental return variance for asset C. Independent noise sources drive the three channels, which add to leading order.

4.1 Case 3: Heterogeneous risk aversion

With $\gamma_H \neq \gamma_L$, per-dollar demand for C differs across types, making the wealth share ω_t a payoff-relevant state variable. Define the heterogeneity parameter

$$\varepsilon = \frac{\gamma_H - \gamma_L}{\gamma_H + \gamma_L}, \qquad \bar{\gamma} = \frac{\gamma_H + \gamma_L}{2}, \tag{14}$$

so that $\gamma_H = \bar{\gamma}(1+\varepsilon)$ and $\gamma_L = \bar{\gamma}(1-\varepsilon)$.

Aggregate risk tolerance. The wealth-weighted average inverse risk aversion is

$$\mathcal{R}(\omega) = \frac{\omega}{\gamma_H} + \frac{1-\omega}{\gamma_L} = \frac{1}{\bar{\gamma}}\big[1 - \varepsilon(2\omega - 1)\big] + O(\varepsilon^2). \tag{15}$$

Market clearing for C gives $\mu_C^{eq} - r = w_C(\sigma_C^{eq})^2/\mathcal{R}(\omega)$, which now depends on ω.

Price-dividend ratio. Since C's risk premium depends on ω, the price-dividend ratio $Q_C(s,\omega)$ is no longer constant. Define $\eta_\omega = \partial \log Q_C/\partial\omega$, the semi-elasticity. The return on C acquires a component from the wealth-share dynamics:

$$\mathrm{d}R_C^{\text{diff}} = \sigma_C\,\mathrm{d}W_u + \eta_\omega\,\mathrm{d}\omega^{\text{diff}}. \tag{16}$$

Wealth-share diffusion. Portfolio weight differences at leading order are $\Delta\pi_C = -2\varepsilon w_C + O(\varepsilon^2)$ and $\Delta\pi_M = -2\varepsilon w_M(s) + O(\varepsilon^2)$. The diffusion of ω is

$$\mathrm{d}\omega^{\text{diff}} = -2\varepsilon\,\omega(1-\omega)\big[w_M(s)\sigma_M(s)\,\mathrm{d}W_F + w_C\sigma_C\,\mathrm{d}W_u\big] + O(\varepsilon^2). \tag{17}$$

Loading on W_u. Since $\eta_\omega = O(\varepsilon)$ and the W_u-loading of $\mathrm{d}\omega$ is $O(\varepsilon)$, the additional return variance from the ω-channel enters at $O(\varepsilon^2)$.

Proposition 3 (Invariance breaks at order ε^2)**.** *In the Duffie-Epstein economy with $\gamma_H \neq \gamma_L$ and two-state regime switching:*

(a) The instantaneous return variance of C is

$$\frac{\mathrm{Var}(\mathrm{d}R_C)}{\mathrm{d}t} = \sigma_C^2 + \Delta_\gamma(s,\omega) + O(\varepsilon^3), \tag{18}$$

where $\Delta_\gamma = O(\varepsilon^2) \geq 0$.

(b) The deviation is proportional to $\omega(1-\omega)$: it vanishes when the wealth distribution is degenerate and peaks at $\omega = 1/2$.

(c) For $\gamma_H = 5.5$, $\gamma_L = 4.5$ (so $\varepsilon = 0.1$), the deviation satisfies $\Delta_\gamma/\sigma_C^2 < 0.03\%$ under standard parameters.

Proof. See Appendix A.1. □

A positive shock $\mathrm{d}W_u$ raises C's return, redistributing wealth toward the less risk-averse type (who holds more C). This lowers aggregate risk aversion, reduces C's discount rate, and raises Q_C. The price response amplifies the fundamental shock. The effect is second-order because it requires both different portfolio weights ($O(\varepsilon)$) and price sensitivity to ω ($O(\varepsilon)$).

4.2 Case 4: Institutional mandates

Type L faces the mandate penalty (7). The first-order condition for π_C^L becomes

$$\pi_C^L = \frac{\mu_C^{eq} - r + \chi\tau_t}{\gamma_L(\sigma_C^{eq})^2 + \chi} = \omega_L \cdot \frac{\mu_C^{eq} - r}{\gamma_L(\sigma_C^{eq})^2} + (1-\omega_L)\tau_t, \tag{19}$$

where $\omega_L = \gamma_L\sigma_C^2/(\gamma_L\sigma_C^2 + \chi)$ is the return-signal weight. The mandated investor's demand for C is a weighted average of the myopic demand and the mandate target.

Market clearing gives a risk premium that depends on τ_t:

$$\mu_C^{eq} - r = \frac{[w_C - (1-\omega)(1-\omega_L)\tau_t](\sigma_C^{eq})^2}{\mathcal{R}^{\text{eff}}(\omega)}, \tag{20}$$

where $\mathcal{R}^{\text{eff}}(\omega) = \omega/\gamma_H + (1-\omega)\omega_L/\gamma_L$ is the effective aggregate risk tolerance. The stochastic τ_t makes the risk premium stochastic through a channel orthogonal to W_F, W_u, and the regime process.

The price-dividend ratio $Q_C(s, \omega, \tau)$ responds to τ with semi-elasticity $\eta_\tau = \partial \log Q_C / \partial \tau$. Using the Gordon growth approximation:

$$\eta_\tau = \frac{(1-\omega)(1-\omega_L)\sigma_C^2}{\delta \cdot \mathcal{R}^{\text{eff}}(\omega)}, \tag{21}$$

where δ is the effective discount rate for C.

Proposition 4 (Mandate-driven variance)**.** *With the mandate target* (6)*:*

(a) The mandate contribution to C's return variance is

$$V_M(s, \omega) = \eta_\tau^2 \, \sigma_\tau(s)^2. \tag{22}$$

(b) $\partial V_M / \partial \sigma_\tau > 0$*: more volatile mandates increase non-fundamental variance.*

(c) $\partial V_M / \partial \chi > 0$*: stricter mandates increase non-fundamental variance.*

(d) $\partial V_M / \partial \omega < 0$*: more unconstrained-investor wealth reduces mandate-driven variance.*

Proof. Part (a): the return on C loads on W_τ through $\eta_\tau \sigma_\tau(s)\, \mathrm{d}W_\tau$. Since W_τ is independent of all other noise sources, the contribution is $\eta_\tau^2 \sigma_\tau^2$. Parts (b)–(d): see Appendix A.2. □

Proposition 5 (Countercyclicality condition)**.** *Define* $\Lambda = V_F(B)/V_F(G)$ *(the ratio of systematic variances),* $\Psi = V_M(B)/V_M(G)$ *(the mandate amplification ratio), and* $\alpha = \sigma_C^2 / V_M(G)$*. The unpriced variance share* S_U *satisfies* $S_U(B) > S_U(G)$ *if and only if*

$$\Psi > \Lambda(1+\alpha) - \alpha. \tag{23}$$

Proof. See Appendix A.3. □

The condition requires mandate-shock volatility to increase across regimes by more than systematic volatility, adjusted for the ratio of fundamental to mandate variance. With standard parameters, the condition fails at baseline: procyclical S_U is the default prediction.

Countercyclical S_U requires severe crisis-induced mandate amplification ($\sigma_\tau(B)/\sigma_\tau(G)$ much larger than $\sigma_M(B)/\sigma_M(G)$).

4.3 Case 5: Stochastic volatility

With v_t following the CIR process (8), the price-dividend ratio $Q_C(v_t)$ depends on the level of variance. The effective discount rate for C is $\delta_C(v_t) = \delta_0+\delta_1 v_t$, where $\delta_1 = \gamma w_C+\frac{1}{2}\gamma(\gamma+1)w_C^2$. Since $\delta_1 > 0$, Q_C is decreasing in v_t: higher fundamental volatility raises the discount rate and lowers the price.

The return on C acquires a volatility-driven component:

$$\mathrm{d}R_C = (\cdots)\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_u + \frac{Q_C'(v_t)}{Q_C(v_t)}\sigma_v\sqrt{v_t}\,\mathrm{d}W_v. \tag{24}$$

Proposition 6 (Stochastic volatility and return variance)**.** *With common γ and CIR volatility:*

(a) The portfolio weight on C is $\pi_C = w_C(s)$, independent of v_t.

(b) The instantaneous return variance decomposes as

$$\frac{\mathrm{Var}(\,\mathrm{d}R_C)}{\mathrm{d}t} = v_t + V_v(v_t), \qquad V_v(v_t) = \left(\frac{Q_C'(v_t)}{Q_C(v_t)}\right)^2 \sigma_v^2 v_t > 0. \tag{25}$$

(c) To first order, $V_v \approx \langle B\rangle^2\sigma_v^2 v_t$, where $\langle B\rangle = \delta_1/(\delta_C+\kappa_v)$ is the duration-weighted Riccati coefficient.

Proof. Part (a): the risk premium $\mu_C - r = \gamma w_C v_t$ scales linearly in v_t, as does the denominator of the demand function γv_t. The ratio is w_C, independent of v_t. Parts (b)–(c): see Appendix A.4. □

Even though v_t is stochastic, the optimal weight on C does not depend on v_t. Disagreement about v_t is therefore irrelevant for portfolio weights. The non-fundamental variance arises entirely through the price level: Q_C responds to v_t, creating capital gains and losses unrelated to dividend news.

Volatility-disagreement stabilization. When Type H observes v_t exactly and Type L observes it with attenuation $\alpha \in [0,1]$ (so $\hat{v}_t^L = \alpha v_t+(1-\alpha)\bar{v}$), the market's effective variance is $v_t^{\mathrm{mkt}} = [\omega+(1-\omega)\alpha]v_t + (1-\omega)(1-\alpha)\bar{v}$.

Proposition 7 (Vol-disagreement stabilization)**.** *With heterogeneous volatility perception (Type L has attenuation $\alpha < 1$):*

$$V_v^{het} = \big[\omega + (1-\omega)\alpha\big]^2 V_v^{hom} < V_v^{hom}. \tag{26}$$

The reduction factor lies in $[\omega^2, 1)$: it decreases in the mass of inattentive investors and in the degree of attenuation.

Proof. The chain rule gives $\partial Q_C/\partial v_t|_{\text{het}} = Q_C'(v_t^{\text{mkt}}) \cdot [\omega + (1-\omega)\alpha]$, and $V_v^{\text{het}} \propto (\partial Q_C/\partial v_t)^2$. □

The stabilization is the opposite of what the standard disagreement literature predicts. Disagreement about the level of a systematic factor typically amplifies price volatility. Here, disagreement about v_t dampens the price response to volatility shocks because the inattentive investors anchor on the long-run mean, providing a stabilizing force.

4.4 Case 6: The full decomposition

With all state variables active, the return on C loads on four independent Brownian motions:

$$\mathrm{d}R_C = (\cdots)\,\mathrm{d}t + \sigma_{R,u}\,\mathrm{d}W_u + \sigma_{R,F}\,\mathrm{d}W_F + \sigma_{R,v}\,\mathrm{d}W_v + \sigma_{R,\tau}\,\mathrm{d}W_\tau, \tag{27}$$

where $\sigma_{R,u}$ contains the fundamental plus the ω-feedback through W_u, $\sigma_{R,F}$ captures the ω-feedback through W_F, $\sigma_{R,v}$ is the volatility channel, and $\sigma_{R,\tau}$ is the mandate channel.

Theorem 8 (Full variance decomposition)**.** *In the full model with regime switching, heterogeneous γ, mandates, and stochastic volatility:*

(a) The instantaneous return variance of C decomposes as

$$\frac{\mathrm{Var}(\mathrm{d}R_C)}{\mathrm{d}t} = \underbrace{v_t}_{\text{fundamental}} + \underbrace{\Delta_\gamma(\omega, s, v)}_{O(\varepsilon^2)} + \underbrace{V_M(\tau, s)}_{O(\sigma_\tau^2)} + \underbrace{V_v(v, s)}_{O(\sigma_v^2)} + O(\varepsilon^3). \tag{28}$$

(b) The three non-fundamental channels are additive to leading order: V_M is driven by W_τ, V_v by W_v, and Δ_γ by W_F and W_u through ω, with no first-order cross terms.

(c) Vol-disagreement reduces V_v by factor $[\omega + (1-\omega)\alpha]^2 < 1$.

(d) Setting $\sigma_v = 0$ recovers Cases 1–4. Setting $\varepsilon = 0$ eliminates Δ_γ. Setting $\sigma_\tau = 0$ eliminates V_M. Setting all three to zero recovers exact invariance.

(e) *The magnitude ordering depends on the characteristic: $V_M > V_v \gg \Delta_\gamma$ for high-mandate characteristics; $V_v > V_M \gg \Delta_\gamma$ for low-mandate characteristics; Δ_γ is always the smallest channel.*

Proof. Part (a): collect the squared loadings on each independent Brownian motion in (27). Part (b): mutual independence of W_u, W_F, W_v, W_τ eliminates cross terms. Part (c): Proposition 7. Part (d): direct substitution. Part (e): established by calibration in Section 5. □

5 Calibration

5.1 Parameters

Table 2 reports the baseline parameter values.

Table 2: Baseline parameters

Parameter	Symbol	Value	Source
Mean risk aversion	$\bar{\gamma}$	5	Standard macro-finance
EIS	ψ	1.5	Bansal and Yaron [2004]
Risk-aversion heterogeneity	ε	0.1	$\gamma_H = 5.5$, $\gamma_L = 4.5$
Regime: $G \to B$ intensity	λ_{GB}	0.05	20-year average expansion
Regime: $B \to G$ intensity	λ_{BG}	0.15	7-year average recession
Systematic dividend drift (G/B)	μ_M	0.08 / 0.02	
Systematic dividend vol (G/B)	σ_M	0.15 / 0.25	
Characteristic dividend drift	μ_D^C	0.05	
Mean characteristic variance	$\bar{v}$	0.04	$\sigma_C = 0.20$
Vol mean-reversion speed	κ_v	0.5	Half-life $\approx$ 1.4 years
Vol-of-vol	σ_v	0.10	
Mandate mean-reversion	κ	0.1	
Mandate target	$\bar{\pi}_C$	0.35	Overweights C vs. $w_C \approx 0.30$
Mandate-shock vol (G/B)	σ_τ	0.05 / 0.10	
Vol attenuation (Type L)	α	0.5	
Wealth share	ω	0.50	Benchmark
Characteristic weight	w_C	0.30	

The perturbation parameter $\varepsilon = 0.1$ is small enough for the expansion to be accurate ($O(\varepsilon^3) = O(0.001)$). The Feller condition $2\kappa_v\bar{v} = 0.04 > \sigma_v^2 = 0.01$ holds.

5.2 Decomposition by regime

Table 3 reports the variance decomposition at the baseline.

The stochastic-volatility channel dominates mandates at baseline. The ratio $V_v/V_M \approx 5$ in expansion and ≈ 8 in recession. The vol channel benefits from the persistent CIR dynamics:

Table 3: Variance decomposition: baseline calibration

Channel	Expansion (G)		Recession (B)	
	Level	% of $\bar{v}$	Level	% of $\bar{v}$
Fundamental v_t	0.0400	100%	0.0400	100%
Mandates V_M	0.00033	0.82%	0.00035	0.87%
Stochastic vol V_v^{het}	0.00165	4.1%	0.00283	7.1%
Wealth redistribution Δ_γ	≈ 0	0.006%	≈ 0	0.027%
Total $\text{Var}(\,\text{d}R_C)/\,\text{d}t$	0.04198		0.04318	
Excess variance ratio		5.0%		8.0%

the duration-weighted Riccati coefficient $\langle B \rangle \approx 2.7$ amplifies σ_v considerably, while the mandate price-impact $\eta_\tau \approx 0.36$ is modest with a mandate distortion $|\bar{\pi}_C - w_C| = 0.05$.

Wealth redistribution is negligible. At $\varepsilon = 0.1$, the Δ_γ channel contributes less than 0.03% of fundamental variance, three orders of magnitude below V_v.

The unpriced share S_U is procyclical at baseline. For a long-short portfolio with characteristic-SDF correlation $\rho = 0.3$, $S_U(G) = 0.950$ and $S_U(B) = 0.875$. The increase in systematic variance from $V_F(G) = 0.0225$ to $V_F(B) = 0.0625$ dominates the modest increase in V_v.

5.3 Cross-characteristic variation

The channel ordering depends on the mandate distortion $|\bar{\pi}_C - w_C|$. Table 4 reports the decomposition for different characteristic types, varying the mandate parameters while holding all other parameters at baseline.

Table 4: Channel ordering by characteristic type (expansion)

Characteristic	$\lvert\bar{\pi}_C - w_C\rvert$	$V_M/\bar{v}$	$V_v/\bar{v}$	Dominant
Low mandate (momentum)	0.02	0.1%	4.1%	V_v
Moderate (value)	0.05	0.8%	4.1%	V_v
High mandate (size)	0.10	3.3%	4.1%	$V_v \approx V_M$
Very high (index-tracking)	0.20	13.1%	4.1%	V_M

The stochastic-volatility channel is parameter-independent across characteristics (it depends on σ_v, κ_v, and δ_1, which are characteristic-specific only through w_C). The mandate channel scales as the square of the mandate distortion. Mandate-driven variance overtakes stochastic-volatility variance when $|\bar{\pi}_C - w_C| \geq 0.15$.

5.4 Countercyclicality assessment

The countercyclicality condition (23) requires $\Psi > \Lambda(1+\alpha) - \alpha$. At baseline: $\Lambda = V_F(B)/V_F(G) = 2.78$, $\alpha = \sigma_C^2/V_M(G) = 122$, and $\Psi = V_M(B)/V_M(G) = 1.07$. The condition requires $\Psi > 220$, which fails by two orders of magnitude. Mandate-shock volatility must increase far more aggressively across regimes than systematic volatility for countercyclical S_U to obtain.

Three features explain the difficulty. First, Λ is large: systematic volatility roughly doubles in recessions. Second, α is large: fundamental variance dominates mandate variance by a factor of 122, so the mandate channel must work very hard to move the ratio. Third, the mandate "sign flip" attenuates countercyclicality: in recession, w_C rises above $\bar{\pi}_C$ because P_M falls, reversing the direction of the mandate distortion.

Countercyclical S_U requires mandate distortions of $|\bar{\pi}_C - w_C| > 0.15$, crisis amplification of $\sigma_\tau(B)/\sigma_\tau(G) > 5$, and mandates that overweight C even after market drawdowns raise w_C. This combination is plausible only for characteristics with extreme institutional demand pressure during crises.

6 Testable Predictions

The model generates eight predictions, organized by the channel that produces them.

6.1 From the invariance

1. **Belief-heterogeneity measures have zero explanatory power for S_U.** Analyst forecast dispersion, options-implied disagreement, and retail-institutional sentiment divergence should not predict the unpriced variance share of characteristic-sorted portfolios after controlling for the characteristic-SDF correlation ρ and institutional demand measures. This zero-coefficient prediction distinguishes the model from behavioral frameworks where belief distortions generate non-fundamental variance.

   Test: Regress S_U (from the Chernov et al. decomposition) on ρ and proxies for mandate distortion (Φ), then add belief-heterogeneity measures. The model predicts zero coefficients on the belief variables.

2. **EIS heterogeneity is irrelevant for S_U.** Cross-sectional variation in investor horizons or intertemporal substitution elasticities (proxied by investor turnover rates or fund redemption frequencies) should not predict S_U after controlling for ρ and Φ.

6.2 From the mandate channel

3. **Mandate segmentation predicts S_U.** Characteristics that define sharp institutional clienteles (size, value/growth) have higher unpriced variance shares than characteristics that cut across mandate boundaries (profitability, investment), after controlling for ρ.

   Test: Use 13F filings to measure the cross-sectional dispersion of portfolio tilts across fund-style categories for each characteristic. Higher dispersion proxies for higher Φ. Regress S_U on ρ and the mandate-dispersion measure.

4. **Mandate-flow volatility predicts S_U.** Characteristics whose associated mandate categories experience more volatile fund flows have higher unpriced variance shares.

   Test: Measure quarterly flow volatility by Morningstar style category. Characteristics associated with high-flow-volatility categories should have higher S_U.

5. **Countercyclical S_U requires extreme mandate amplification.** Procyclical S_U is the baseline for most characteristics. Countercyclical S_U obtains only for characteristics where mandate-shock volatility increases by a factor exceeding $\Lambda(1+\alpha)-\alpha$, where Λ is the ratio of systematic variances.

   Test: Sort characteristics by institutional mandate intensity. High-mandate characteristics should exhibit less procyclical (or countercyclical) S_U in crisis periods, while low-mandate characteristics should exhibit strongly procyclical S_U.

6.3 From the stochastic-volatility channel

6. **Vol-of-vol predicts S_U.** Characteristics with higher volatility of idiosyncratic volatility (higher σ_v) have higher unpriced variance shares, holding mandate parameters fixed.

   Test: Estimate σ_v from high-frequency data for each characteristic portfolio. Regress S_U on σ_v, ρ, and Φ.

7. **Vol-disagreement stabilizes.** In the cross-section, characteristics for which institutional and retail investors have more dispersed volatility estimates should have *lower* excess variance. Disagreement about v_t dampens the price response to volatility shocks.

   Test: Construct a vol-disagreement measure from the dispersion of implied volatilities across option maturities or from the spread between realized and option-implied volatility. The model predicts a negative coefficient on this measure in a regression of S_U on vol-disagreement, controlling for σ_v.

6.4 From the cross-characteristic ordering

8. **The channel ordering is observable.** Decomposing S_U into a component correlated with vol-of-vol measures and a component correlated with mandate-flow measures should reveal that mandate-correlated excess variance dominates for size and value, while vol-correlated excess variance dominates for momentum and short-term reversal.

   Test: Regress time-series variation in S_U for each characteristic on both vol-of-vol and mandate-flow-volatility measures. The relative R^2 contributions should vary across characteristics in the pattern predicted by Table 4.

6.5 Summary of comparative statics

Table 5 collects the model's predictions.

Table 5: Comparative statics for S_U

Variable	Effect on S_U	Source
ρ (characteristic-SDF correlation)	Decreases	Definition
Φ (mandate segmentation intensity)	Increases	Proposition 4
σ_τ (mandate-flow volatility)	Increases	Proposition 4(b)
χ (mandate strictness)	Increases	Proposition 4(c)
ω (unconstrained-investor wealth)	Decreases	Proposition 4(d)
σ_v (vol-of-vol)	Increases	Proposition 6
Vol-disagreement	Decreases	Proposition 7
Belief heterogeneity (any form)	No effect	Propositions 1–2
EIS heterogeneity	No effect	Proposition 2
ε (risk-aversion heterogeneity)	Increases ($O(\varepsilon^2)$)	Proposition 3

7 Conclusion

A continuous-time economy with Epstein-Zin preferences, regime switching, institutional mandates, and stochastic volatility delivers a dichotomy for the unpriced variance of characteristic-sorted portfolios. Beliefs about systematic factors cannot change it: the invariance holds exactly under common risk aversion and approximately (deviation below 1%) under heterogeneous risk aversion. Two first-order channels break the invariance: institutional mandate shocks inject non-fundamental variance proportional to mandate-flow volatility, and stochastic characteristic-specific volatility creates price variation through the discount-rate channel, with disagreement about volatility providing a stabilizing force. The relative importance of these channels varies across characteristics, generating testable predictions that link the cross-section of unpriced variance to observable institutional demand patterns and volatility dynamics.

References

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

Suleyman Basak and Domenico Cuoco. An equilibrium model with restricted stock market participation. *Review of Financial Studies*, 11(2):309–341, 1998.

Mikhail Chernov, Magnus Dahlquist, and Lars Lochstoer. Unpriced risks: Rethinking cross-sectional asset pricing. *NBER Working Paper 34009*, 2025.

Darrell Duffie and Larry G. Epstein. Stochastic differential utility. *Econometrica*, 60(2): 353–394, 1992.

Darrell Duffie, Nicolae Gârleanu, and Lasse Heje Pedersen. Over-the-counter markets. *Econometrica*, 73(6):1815–1847, 2005.

Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.

Ralph S.J. Koijen and Motohiro Yogo. A demand system approach to asset pricing. *Journal of Political Economy*, 127(4):1475–1515, 2019.

Costis Skiadas. *Asset Pricing Theory*. Princeton University Press, 2009.

Dimitri Vayanos and Jean-Luc Vila. A preferred-habitat model of the term structure of interest rates. *Econometrica*, 89(1):77–112, 2021.

A Proofs

A.1 Proof of Proposition 3 (Case 3: wealth redistribution)

Proof. **Step 1: Portfolio weight differences.** With $\gamma_H = \bar{\gamma}(1+\varepsilon)$ and $\gamma_L = \bar{\gamma}(1-\varepsilon)$:

$$\frac{1}{\gamma_H} - \frac{1}{\gamma_L} = \frac{1}{\bar{\gamma}}\left(\frac{1}{1+\varepsilon} - \frac{1}{1-\varepsilon}\right) = -\frac{2\varepsilon}{\bar{\gamma}(1-\varepsilon^2)}.$$

From the myopic demand $\pi_C^j = (\mu_C - r)/(\gamma_j \sigma_C^2)$ and market clearing $\mu_C - r = w_C\sigma_C^2/\mathcal{R}(\omega)$:

$$\Delta\pi_C = \pi_C^H - \pi_C^L = \frac{w_C}{\mathcal{R}(\omega)}\left(\frac{1}{\gamma_H} - \frac{1}{\gamma_L}\right) = -2\varepsilon w_C + O(\varepsilon^2).$$

Similarly, $\Delta\pi_M = -2\varepsilon w_M(s) + O(\varepsilon^2)$.

Step 2: Wealth-share diffusion. The diffusion of ω_t is

$$\mathrm{d}\omega^{\text{diff}} = \omega(1-\omega)\big[\Delta\pi_M\sigma_M(s)\,\mathrm{d}W_F + \Delta\pi_C\sigma_C\,\mathrm{d}W_u\big].$$

Substituting: $\mathrm{d}\omega^{\text{diff}} = -2\varepsilon\omega(1-\omega)[w_M\sigma_M\,\mathrm{d}W_F + w_C\sigma_C\,\mathrm{d}W_u] + O(\varepsilon^2)$.

Step 3: Semi-elasticity. The P/D ratio satisfies $Q_C \approx 1/\delta$ where $\delta = r + w_C\sigma_C^2/\mathcal{R}(\omega) - \mu_D^C$. Define $\phi = \partial\delta/\partial\omega|_{\varepsilon=0+}$. Then $\eta_\omega = -\varepsilon\phi/\delta_0 + O(\varepsilon^2)$.

Step 4: Return variance. The return on C loads on W_u with coefficient

$$b_u = \sigma_C + \eta_\omega\cdot(-2\varepsilon)\omega(1-\omega)w_C\sigma_C = \sigma_C\left[1 + \frac{2\varepsilon^2\phi}{\delta_0}\omega(1-\omega)w_C\right] + O(\varepsilon^3).$$

The loading on W_F is $b_F = (2\varepsilon^2\phi/\delta_0)\omega(1-\omega)w_M\sigma_M + O(\varepsilon^3)$. Then:

$$\frac{\mathrm{Var}(\,\mathrm{d}R_C)}{\mathrm{d}t} = b_u^2 + b_F^2 = \sigma_C^2 + \frac{4\varepsilon^2\phi\sigma_C^2 w_C}{\delta_0}\omega(1-\omega) + O(\varepsilon^3).$$

The deviation $\Delta_\gamma = (4\varepsilon^2\phi\sigma_C^2 w_C/\delta_0)\omega(1-\omega) \geq 0$ under $\phi > 0$ (the risk-premium effect of higher aggregate risk aversion dominates the risk-free-rate effect). The maximum occurs at $\omega = 1/2$, establishing parts (a) and (b).

Step 5: Quantification. With $\varepsilon = 0.1$, $\phi w_C/\delta_0 \approx 1$, and $\omega(1-\omega) = 0.25$:

$$\frac{\Delta_\gamma}{\sigma_C^2} = 4(0.01)(1)(0.25) = 0.01 = 1\%.$$

The calibration in Section 5 yields 0.006%–0.027%, well below 1%. □

A.2 Proof of Proposition 4 (Case 4: mandates)

Proof. **Part (a).** From (19), market clearing gives the risk premium (20). The P/D ratio $Q_C(s, \omega, \tau)$ responds to τ through η_τ (equation (21)). The return on C loads on W_τ with coefficient $\eta_\tau \sigma_\tau(s)$. Since $W_\tau \perp W_F, W_u, W_v$, the variance contribution is $V_M = \eta_\tau^2 \sigma_\tau(s)^2$.

Part (b). $V_M = \eta_\tau^2 \sigma_\tau^2$ with $\eta_\tau > 0$. Then $\partial V_M / \partial \sigma_\tau = 2\eta_\tau^2 \sigma_\tau > 0$.

Part (c). Increasing χ reduces $\omega_L = \gamma_L \sigma_C^2 / (\gamma_L \sigma_C^2 + \chi)$:

$$\frac{\partial(1-\omega_L)}{\partial \chi} = \frac{\gamma_L \sigma_C^2}{(\gamma_L \sigma_C^2 + \chi)^2} > 0.$$

This increases $(1-\omega_L)$ in η_τ. Simultaneously, $\mathcal{R}^{\text{eff}}$ decreases (since ω_L falls), so $\Gamma_C^{\text{eff}} = 1/\mathcal{R}^{\text{eff}}$ increases. Both effects raise η_τ, so $\partial V_M / \partial \chi > 0$.

Part (d). Define $g(\omega) = (1-\omega)/\mathcal{R}^{\text{eff}}(\omega)$. Then $V_M \propto g(\omega)^2$.

$$g(\omega) = \frac{1-\omega}{\omega/\gamma_H + (1-\omega)\omega_L/\gamma_L}.$$

Differentiating the numerator by ω gives -1. Differentiating the denominator gives $1/\gamma_H - \omega_L/\gamma_L$. By the quotient rule, the numerator of g' is

$$-\frac{\omega}{\gamma_H} - \frac{(1-\omega)\omega_L}{\gamma_L} - (1-\omega)\left(\frac{1}{\gamma_H} - \frac{\omega_L}{\gamma_L}\right) = -\frac{1}{\gamma_H} < 0.$$

So $g' < 0$ and $\partial V_M / \partial \omega < 0$. □

A.3 Proof of Proposition 5 (Countercyclicality)

Proof. For the long-short portfolio with characteristic-SDF correlation ρ:

$$S_U(s) = \frac{(1-\rho^2)[\sigma_C^2 + V_M(s)]}{\rho^2 V_F(s) + (1-\rho^2)[\sigma_C^2 + V_M(s)]}.$$

Write $S_U = 1/(1+R)$ with $R(s) = \rho^2 V_F(s)/\{(1-\rho^2)[\sigma_C^2 + V_M(s)]\}$. Then $S_U(B) > S_U(G) \iff R(B) < R(G) \iff$

$$\frac{V_F(B)}{\sigma_C^2 + V_M(B)} < \frac{V_F(G)}{\sigma_C^2 + V_M(G)}.$$

Cross-multiplying and rearranging:

$$V_F(B)[\sigma_C^2 + V_M(G)] < V_F(G)[\sigma_C^2 + V_M(B)].$$

Dividing both sides by $V_F(G)V_M(G)$ and using $\Lambda = V_F(B)/V_F(G)$, $\Psi = V_M(B)/V_M(G)$, $\alpha = \sigma_C^2/V_M(G)$:

$$\Lambda(\alpha+1) < \alpha + \Psi, \qquad \text{i.e.,} \qquad \Psi > \Lambda(1+\alpha) - \alpha. \qquad \square$$

A.4 Proof of Proposition 6 (Stochastic volatility)

Proof. **Part (a).** The equilibrium risk premium is $\mu_C - r = \gamma w_C v_t$. The demand for C is

$$\pi_C = \frac{\mu_C - r}{\gamma v_t} = \frac{\gamma w_C v_t}{\gamma v_t} = w_C,$$

independent of v_t.

Part (b). With $P_C = Q_C(v_t)D_C$, Ito's lemma gives:

$$\mathrm{d}R_C = \left[\frac{1}{Q_C} + \mu_D^C + \frac{Q_C'}{Q_C}\kappa_v(\bar{v} - v_t) + \frac{Q_C''}{2Q_C}\sigma_v^2 v_t\right]\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_u + \frac{Q_C'}{Q_C}\sigma_v\sqrt{v_t}\,\mathrm{d}W_v.$$

The cross-term $\mathrm{d}Q_C \cdot \mathrm{d}D_C = 0$ because $\langle W_u, W_v\rangle = 0$. Since $W_u \perp W_v$:

$$\frac{\mathrm{Var}(\mathrm{d}R_C)}{\mathrm{d}t} = v_t + \left(\frac{Q_C'}{Q_C}\right)^2 \sigma_v^2 v_t = v_t + V_v(v_t).$$

Since $Q_C' < 0$ (higher variance raises the discount rate), $V_v > 0$.

Part (c). For the exponential-affine form $Q_C = \int_0^\infty e^{\alpha(\tau)+\beta(\tau)v_t}\,\mathrm{d}\tau$, the cash-flow-weighted sensitivity is $Q_C'/Q_C \approx \langle B\rangle$ where

$$\langle B\rangle = \frac{\delta_1}{\kappa_v}\left[1 - \frac{\delta_C}{\delta_C + \kappa_v}\right] = \frac{\delta_1}{\delta_C + \kappa_v}.$$

Therefore $V_v \approx \langle B\rangle^2\sigma_v^2 v_t$. $\square$

B Extensions of the Invariance

B.1 Ambiguity aversion

Consider the smooth ambiguity model of Klibanoff et al. [2005]. Agent j evaluates asset C using

$$U_j = \mathbb{E}_\pi\big[\phi_j(\mathbb{E}_\mu[u(W^j)])\big],$$

where π is a prior over models μ, and ϕ_j is a concave transformation capturing ambiguity aversion. If ambiguity concerns the systematic factor F (the set of models μ differs in beliefs about F but not about u), then the inner expectation $\mathbb{E}_\mu[u(W^j)]$ depends on μ only through the F-related component of wealth. The C-component separates:

$$\mathbb{E}_\mu[u(W^j)] = \mathbb{E}_\mu\big[u(W_M^j + W_C^j)\big],$$

where W_C^j is the wealth from the C-position, independent of μ. The first-order condition for π_C^j does not depend on ϕ_j or on the prior π over F-models.

Proposition 9. *For any smooth, strictly increasing, concave ambiguity function ϕ, the invariance of $\mathrm{Var}(\,\mathrm{d}R_C)/\,\mathrm{d}t$ to beliefs about systematic factors holds.*

B.2 Invariance in the noisy rational expectations equilibrium

Suppose investors receive private signals about u: agent j observes $s_j = u + \eta_j$, $\eta_j \sim N(0, \sigma_\eta^2)$. In the Grossman and Stiglitz [1980] tradition, the equilibrium price of C partially reveals the aggregate signal.

In this richer information environment, the return on C reflects both fundamental risk (u) and noise-trader demand. Beliefs about F remain irrelevant for C's pricing because u and F are independent. The REE price of C aggregates information about u from private signals and the price itself, but this aggregation does not depend on beliefs about F.

Proposition 10. *In the Grossman-Stiglitz NREE with private signals about u and heterogeneous beliefs about F, the equilibrium return variance of C does not depend on the distribution of F-beliefs.*

The proof follows from the block-diagonal structure of the signal-extraction problem: signals about u are independent of signals about F, so the REE price of C does not load on F-beliefs.

B.3 Invariance under search frictions

In an OTC market where investors trade C through bilateral search (Duffie et al. 2005), the surplus from a trade in C between agents i and j is

$$S_{ij} = (v_i^C - v_j^C),$$

where v_i^C is agent i's valuation of C. Since C's payoff is independent of F, each agent's valuation v_i^C depends on beliefs about u and risk aversion, not on beliefs about F. Search frictions slow down trading but do not change the set of valuations. The steady-state price of C and its variance are determined by the distribution of u-valuations and the search technology, independent of F-beliefs.

Proposition 11. *Under bilateral search with Nash bargaining, the equilibrium variance of C's return is invariant to the distribution of beliefs about F.*

C Additional Breaking Channels

Four additional channels break the invariance. Each operates through a distinct mechanism.

Leverage constraints. When investors face a portfolio-level leverage constraint $|\pi_M| + |\pi_C| \leq L$, a binding constraint on the M-position spills over to C: the constraint links demands for M and C even though their payoffs are independent. The invariance breaks as a step function: $\Delta = 0$ when the constraint is slack, $\Delta > 0$ when it binds.

Learning about factor structure. If investors are uncertain about the factor loading β linking C to F and learn Bayesian-ly from return data, the posterior variance of β declines over time. This creates hump-shaped dynamics in S_U: early in the sample, model uncertainty inflates C's variance; as learning accumulates, S_U converges to the benchmark.

Intermediary capital constraints. When financial intermediaries face VaR or risk-budget constraints, a crisis that tightens the constraint on M-positions reduces the intermediary's capacity to absorb C-risk, increasing C's equilibrium return variance. The effect is regime-dependent, amplifying in bad times.

Production economy. If firms choose their factor loadings through investment decisions, beliefs about F can affect the supply of characteristics. A firm that expects high F-returns may invest to increase its F-loading, changing the composition of assets available for characteristic sorting. This supply-side channel operates at $O(\phi^3)$ where ϕ measures the investment-response elasticity.

D Two-Asset Reduction from the Continuum Economy

The two-asset model (assets M and C) can be derived as a reduction from an economy with a continuum of stocks. Suppose stock i pays $D_i = \beta_i F + \varepsilon_i$ where β_i is the factor loading and ε_i is idiosyncratic. A characteristic $c_i = \rho\beta_i + \sqrt{1-\rho^2}\eta_i$ is an imperfect proxy for β_i, where ρ measures the informativeness of the characteristic.

Sorting stocks into a long-short portfolio by c_i and taking the continuum limit yields a portfolio with return

$$r_{\text{LS}} = a_1 F + a_2 u,$$

where u aggregates the idiosyncratic risk that does not diversify away because characteristic sorting tilts toward stocks with correlated ε_i. The coefficients a_1 and a_2 depend on ρ: when $\rho = 1$, the sort perfectly separates high-β from low-β stocks and $a_2 = 0$; when $\rho = 0$, the sort is uninformative and $a_1 = 0$.

The variance decomposition is $\text{Var}(r_{\text{LS}}) = a_1^2 \,\text{Var}(F) + a_2^2 \,\text{Var}(u)$, where the first term is priced (systematic) and the second is unpriced (characteristic-specific). This motivates the two-asset model: asset M has return $r_M = F + \text{noise}$ and asset C has return $r_C = u + \text{noise}$, with the long-short portfolio loading on both according to ρ.