Epstein-Zin Preferences

Lorenzo Naranjo

March 2026

This notebook introduces Epstein-Zin preferences as a way to separate risk aversion from intertemporal substitution in asset-pricing models. It starts with the lognormal benchmark under power utility, then derives how recursive utility changes the stochastic discount factor and risk-premium decomposition. The goal is to build clear intuition for which macro risks are priced, why long-run risk matters, and how these equations connect to the empirical asset-pricing puzzles discussed in class.

Lognormal Consumption

As a benchmark before introducing Epstein-Zin recursion, start with the standard power-utility model and lognormal consumption growth. This baseline pins down how interest rates and risk premia relate to marginal utility when risk aversion and intertemporal substitution are tied together. Let lowercase letters denote logs, so $c_t = \ln(C_t)$ and $\Delta c_{t+1} = c_{t+1} - c_t$. Assume

$$\Delta c_{t+1} \sim \mathcal{N}\left(\mathsf{E}_t(\Delta c_{t+1}), \sigma_t^2(\Delta c_{t+1})\right)$$

and preferences are CRRA (power utility):

$$u(C) = \begin{cases} \frac{C^{1-\gamma}}{1-\gamma}, & \text{if } \gamma \geq 0, \gamma \neq 1 \\ \ln(C), & \text{if } \gamma = 1 \end{cases}$$

so that

$$u'(C) = C^{-\gamma}.$$

We can then write $m_{t+1} = \ln(M_{t+1})$ as:

$$m_{t+1} = -\delta - \gamma \Delta c_{t+1}$$

where $\delta = -\ln(\beta)$. Thus, $m_{t+1} \sim \mathcal{N}(\mathrm{E}_t(m_{t+1}), \sigma_t^2(m_{t+1}))$ where

$$\begin{aligned} \mathrm{E}_t(m_{t+1}) &= -\delta - \gamma\, \mathrm{E}_t(\Delta c_{t+1}), \\ \sigma_t^2(m_{t+1}) &= \gamma^2 \sigma_t^2(\Delta c_{t+1}). \end{aligned}$$

The Risk-Free Rate. We can compute

$$\begin{aligned} \mathrm{E}_t(M_{t+1}) &= \exp\left(\mathrm{E}_t(m_{t+1}) + \frac{1}{2}\sigma_t^2(m_{t+1})\right) \\ &= \exp\left(-\delta - \gamma\, \mathrm{E}_t(\Delta c_{t+1}) + \frac{1}{2}\gamma^2 \sigma_t^2(\Delta c_{t+1})\right) \end{aligned}$$

Therefore,

$$\begin{aligned} r_{f,t+1} = \ln(R_{f,t+1}) &= -\ln(\mathrm{E}_t(M_{t+1})) \\ &= \delta + \gamma\, \mathrm{E}_t(\Delta c_{t+1}) - \frac{1}{2}\gamma^2 \sigma_t^2(\Delta c_{t+1}) \end{aligned}$$

This expression implies three comparative statics for the real risk-free rate. Higher impatience (δ) raises the risk-free rate. Higher expected consumption growth also raises the risk-free rate because if households expect to be richer tomorrow, they save less today, lowering bond demand and raising yields. By contrast, higher consumption-growth uncertainty lowers the risk-free rate through precautionary saving. In the data, the real risk-free rate is low and stable despite positive expected consumption growth—a tension sometimes called the *risk-free rate puzzle*—which Epstein-Zin preferences help address by decoupling the precautionary-saving response from risk aversion.

Equity Premium Puzzle. If $m \sim \mathcal{N}(\bar{m}, \bar{\sigma}^2)$ and $M = e^m$, then $\mathrm{E}(M) = e^{\bar{m}+0.5\bar{\sigma}^2}$ and $\sigma^2(M) = (e^{\bar{\sigma}^2} - 1)e^{2\bar{m}+\bar{\sigma}^2}$. Therefore

$$\frac{\sigma(M)}{\mathrm{E}(M)} = \sqrt{e^{\bar{\sigma}^2} - 1} \approx \bar{\sigma}$$

Using the Hansen-Jagannathan bound, we can write the Sharpe ratio of the market as

$$\left|\frac{\mathrm{E}_t(R_{m,t+1}) - R_{f,t+1}}{\sigma_t(R_{m,t+1})}\right| \leq \gamma \sigma_t(\Delta c_{t+1})$$

In the data, the market Sharpe ratio is around 0.5, whereas the standard deviation of consumption growth is around 0.01, implying an RRA coefficient of at least 50.

Lognormal Returns. Next, assume log returns and the log SDF are jointly conditionally normal, so that:

$$
\begin{aligned}
1 &= \mathsf{E}_t(M_{t+1}R_{i,t+1}) \\
&= \mathsf{E}_t(e^{m_{t+1}+r_{i,t+1}}) \\
&= e^{\mathsf{E}_t(m_{t+1}+r_{i,t+1})+\frac{1}{2}\sigma_t^2(m_{t+1}+r_{i,t+1})}
\end{aligned}
$$

Therefore,

$$
\mathsf{E}_t(m_{t+1}+r_{i,t+1}) + \frac{1}{2}\sigma_t^2(m_{t+1}+r_{i,t+1}) = 0 \tag{1}
$$

Substituting the CRRA log SDF into (1) and expanding $\sigma_t^2(m_{t+1}+r_{i,t+1})$ gives

$$
\mathsf{E}_t\,\Delta c_{t+1} = -\frac{\delta}{\gamma} + \frac{1}{\gamma}\mathsf{E}_t\, r_{i,t+1} + \frac{1}{2}\sigma_t^2\left(\frac{1}{\gamma}r_{i,t+1} - \Delta c_{t+1}\right) \tag{2}
$$

Equation (2) shows that expected consumption growth increases with expected asset returns and with the conditional variance of $\frac{1}{\gamma}r_{i,t+1} - \Delta c_{t+1}$. Expanding that variance implies dependence on return variance, consumption-growth variance, and their covariance.

Equation (1) also implies that

$$
\mathsf{E}_t(r_{i,t+1}) - r_{f,t+1} + \frac{V_{ii,t}}{2} = \gamma V_{ic,t} \tag{3}
$$

where $V_{ii,t} = \sigma_t^2(r_{i,t+1})$ and $V_{ic,t} = \mathrm{Cov}_t(\Delta c_{t+1}, r_{i,t+1})$. Equation (3) says expected excess log return is increasing in an asset's covariance with consumption growth ($V_{ic,t}$). Assets that pay off more when consumption is high are riskier for a representative investor and therefore require higher expected returns.

The term $V_{ii,t}/2$ is a Jensen's inequality correction specific to log returns and is small in practice relative to the covariance term.

Epstein-Zin Preferences

The standard CRRA time-additive model is equivalent to a recursive specification: the continuation value U_t satisfies the Bellman equation

$$U_t^{1-\gamma} = (1-\beta)C_t^{1-\gamma} + \beta\, \mathsf{E}_t(U_{t+1}^{1-\gamma}).$$

In this formulation, the agent's attitude toward risk and their willingness to substitute consumption over time are both governed by the single parameter γ. Specifically, γ is the coefficient of relative risk aversion, while the elasticity of intertemporal substitution (EIS) is constrained to equal $1/\gamma$. This tight link is a restrictive feature of the standard model: an agent who is very risk averse (large γ) is also forced to have a very low willingness to substitute consumption over time, even though these are conceptually distinct aspects of preferences.

The Epstein-Zin model decouples these two parameters by introducing an additional scaling parameter θ into the recursion:

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

where

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

The parameter ψ now independently controls the EIS, so γ and ψ can be chosen freely to match empirical estimates of risk aversion and intertemporal substitution separately. When $\psi = 1/\gamma$ we have $\theta = 1$ and the standard power utility model is recovered as a special case. Intuitively, θ governs the relative weight on intertemporal substitution versus market-return news in the SDF: when $\theta > 1$ (equivalently, $\gamma > 1/\psi$) the agent prefers early resolution of uncertainty, and this tilts the SDF toward the market-return factor.

The Budget Constraint. In any dynamic model of portfolio choice, consumption C_t is derived from the agent's existing wealth W_t. After consuming, the remaining wealth $W_t - C_t$ is invested in a portfolio whose gross return from period t to $t+1$ we denote $R_{w,t+1}$, giving the intertemporal budget constraint

$$W_{t+1} = R_{w,t+1}(W_t - C_t) \tag{4}$$

The return $R_{w,t+1}$ is best understood as the return on a claim to the agent's entire future consumption stream—the so-called *market portfolio* that also appears in the traditional CAPM. Since wealth equals the present discounted value of current and future consumption, we have

$$W_t = \mathrm{E}_t\left(\sum_{j=0}^{\infty} M_{t+j} C_{t+j}\right),$$

where M_{t+j} denotes the multi-period stochastic discount factor from t to $t+j$ (the product of one-period SDFs, with $M_{t+0} \equiv 1$).

The Stochastic Discount Factor. Applying Bellman's optimality condition to the EZ recursion, it can be shown that the stochastic discount factor takes the form

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

This expression makes the separation of preferences transparent. The first factor is an intertemporal substitution term: it depends on consumption growth raised to $-1/\psi$, the inverse of the EIS. The second factor captures the return on the market portfolio and reflects how utility is affected by revisions in future investment opportunities. In the standard power utility case ($\theta = 1$), the second factor vanishes and only consumption growth determines the SDF.

Taking logs, with $\delta = -\ln\beta$ and lowercase letters denoting logs of their uppercase counterparts, the log SDF is

$$m_{t+1} = -\theta\delta - (\theta/\psi)\Delta c_{t+1} - (1-\theta) r_{w,t+1}.$$

Using the general log pricing equation (3), the log risk premium for any asset i satisfies

$$\mathrm{E}_t(r_{i,t+1}) - r_{f,t+1} + \frac{V_{ii,t}}{2} = \frac{\theta}{\psi} V_{ic,t} + (1-\theta) V_{iw,t} \tag{5}$$

where $V_{ic,t} = \mathrm{Cov}_t(r_{i,t+1}, \Delta c_{t+1})$ and $V_{iw,t} = \mathrm{Cov}_t(r_{i,t+1}, r_{w,t+1})$. The left-hand side is the expected excess log return plus a convexity correction. The right-hand side loads on two covariances—with consumption growth and with the market return—with weights that depend on θ and hence on the preference parameters γ and ψ.

A Note on Solving Difference Equations

The log-linearization and innovation decompositions below rely on the solution to a class of forward-looking difference equations. Consider

$$y_t = \rho(x_{t+1} + y_{t+1})$$

where x_t is a random variable known at time t and $|\rho| < 1$. To solve for y_t, we iterate forward. Substituting for y_{t+1}, then for y_{t+2}, and so on yields the telescoping expansion

$$\begin{aligned} y_t &= \rho x_{t+1} + \rho y_{t+1} \\ &= \rho x_{t+1} + \rho^2(x_{t+2} + y_{t+2}) \\ &= \rho x_{t+1} + \rho^2 x_{t+2} + \rho^2 y_{t+2} \\ &= \rho x_{t+1} + \rho^2 x_{t+2} + \rho^3 x_{t+3} + \rho^3 y_{t+3}, \end{aligned}$$

so that

$$y_t = \sum_{j=1}^{n} \rho^j x_{t+j} + \rho^n y_{t+n}.$$

As long as the *transversality condition* $\lim_{n\to\infty} \rho^n y_{t+n} = 0$ holds—which rules out explosive bubble paths in the variable y_t—the remainder term vanishes as $n \to \infty$ and we obtain

$$y_t = \sum_{j=1}^{\infty} \rho^j x_{t+j} \tag{6}$$

This solution holds ex-post, meaning it holds path by path for the actual realizations of x_{t+j}. Since it holds ex-post, it must also hold ex-ante as a conditional expectation:

$$y_t = \mathrm{E}_t \sum_{j=1}^{\infty} \rho^j x_{t+j}.$$

Revisions of Expectations. The solution (6) is determined by the entire sequence of future realizations x_{t+j} for $j \geq 1$. As time passes and new information arrives, agents revise their expectations. At time t the agent expects x_{t+j} to equal $\mathrm{E}_t(x_{t+j})$; a period later, the expectation updates to $\mathsf{E}_{t+1}(x_{t+j})$. We define the *revision* in expectations of x_{t+j} as

$$(\mathsf{E}_{t+1} - \mathsf{E}_t)x_{t+j} = \mathsf{E}_{t+1}(x_{t+j}) - \mathsf{E}_t(x_{t+j}).$$

A key identity follows from the fact that y_t is already known at time t: since $(\mathsf{E}_{t+1} - \mathsf{E}_t)y_t = 0$, we have

$$(\mathsf{E}_{t+1} - \mathsf{E}_t)\sum_{j=1}^{\infty} \rho^j x_{t+j} = 0.$$

This identity—that the discounted sum of all expectation revisions must equal zero when the left-hand variable is predetermined—is what underlies every innovation decomposition in these notebooks.

Log-Linearization of the Budget Constraint

To make the model tractable, Campbell (1993) proposes log-linearizing the budget constraint. Dividing (4) by W_t and taking logs gives

$$\Delta w_{t+1} = r_{w,t+1} + \ln(1 - \exp(c_t - w_t)) \tag{7}$$

where $\Delta w_{t+1} = w_{t+1} - w_t$ denotes log wealth growth. The difficulty is the nonlinear term $\ln(1 - e^{c_t - w_t})$, which depends on the log consumption-wealth ratio $cw_t = c_t - w_t$.

To handle this nonlinearity, we perform a first-order Taylor expansion of $\ln(1 - e^x)$ around its steady-state value $\bar{x}$:

$$\ln(1 - e^x) \approx \ln(1 - e^{\bar{x}}) - \frac{e^{\bar{x}}}{1 - e^{\bar{x}}}(x - \bar{x}) \tag{8}$$

Defining $\rho = 1 - e^{\bar{x}}$, which equals the steady-state fraction of wealth that is saved rather than consumed, this simplifies to

$$\ln(1 - e^x) \approx k + \left(1 - \frac{1}{\rho}\right)x \tag{9}$$

where k can be computed from (8). The parameter ρ is slightly less than one and plays a central discounting role for future terms throughout the model.

Substituting (9) into (7) yields a linear approximation to log wealth growth:

$$\Delta w_{t+1} = k + r_{w,t+1} + \left(1 - \frac{1}{\rho}\right) cw_t$$

where $cw_t = c_t - w_t$. Because $\Delta w_{t+1} = \Delta c_{t+1} - cw_{t+1} + cw_t$ by definition, we can rearrange to write the consumption-wealth ratio as a function of next period's variables:

$$cw_t = \rho(k + r_{w,t+1} - \Delta c_{t+1} + cw_{t+1}).$$

This is a forward-looking difference equation in cw_t. Applying the solution (6), we obtain

$$cw_t = \sum_{j=1}^{\infty} \rho^j (k + r_{w,t+j} - \Delta c_{t+j}) \tag{10}$$

so that the log consumption-wealth ratio at time t equals the discounted sum of all future differences between market returns and consumption growth.

Innovations in Consumption Growth. An important implication of (10) is that, since cw_t is known at time t, the revision in expectations of cw_t between t and $t + 1$ must be zero:

$$(\mathsf{E}_{t+1} - \mathsf{E}_t) \sum_{j=1}^{\infty} \rho^j (r_{w,t+j} - \Delta c_{t+j}) = 0.$$

Expanding and rearranging yields an expression for the unexpected component of consumption growth:

$$\Delta c_{t+1} - \mathsf{E}_t \Delta c_{t+1} = r_{w,t+1} - \mathsf{E}_t r_{w,t+1} + (\mathsf{E}_{t+1} - \mathsf{E}_t) \sum_{j=1}^{\infty} \rho^j (r_{w,t+j+1} - \Delta c_{t+j+1}) \tag{11}$$

Equation (11) reveals three channels through which consumption can grow unexpectedly: an unexpected high return on wealth today, upward revisions in expected future market returns, and upward revisions in expected future consumption growth itself. All three raise current consumption by making the agent feel wealthier via the permanent income logic: good news about any future income stream raises current wealth and thus current spending.

Using the general lognormal pricing equation (1), applied here to the market return under EZ preferences, one can show that expected consumption growth satisfies

$$\mathrm{E}_t\,\Delta c_{t+1} = \mu_t + \psi\,\mathrm{E}_t\,r_{w,t+1} \tag{12}$$

The slope ψ is the EIS: higher expected market returns translate directly into higher expected consumption growth, capturing the intertemporal substitution channel. Here

$$\mu_t = -\delta\psi + \frac{1}{2}\left(\frac{\theta}{\psi}\right)\sigma_t^2\left(\psi r_{w,t+1} - \Delta c_{t+1}\right)$$

captures the precautionary savings motive and terms related to conditional variances. Substituting into (11) to eliminate future consumption growth in favor of future market returns gives

$$\begin{aligned}\Delta c_{t+1} - \mathrm{E}_t\,\Delta c_{t+1} = r_{w,t+1} - \mathrm{E}_t\,r_{w,t+1} \\ + (1-\psi)(\mathrm{E}_{t+1} - \mathrm{E}_t)\sum_{j=1}^{\infty}\rho^j r_{w,t+j+1} \\ - (\mathrm{E}_{t+1} - \mathrm{E}_t)\sum_{j=1}^{\infty}\rho^j \mu_{t+j+1}\end{aligned} \tag{13}$$

Rewriting the Covariance with Consumption Growth. Equation (13) implies that the conditional covariance of any asset return $r_{i,t+1}$ with consumption growth decomposes as

$$\mathrm{Cov}_t(r_{i,t+1}, \Delta c_{t+1}) = V_{iw,t} + (1-\psi)V_{ih,t} - V_{i\sigma,t}$$

where

$$V_{iw,t} = \text{Cov}_t(r_{i,t+1}, r_{w,t+1})$$
$$V_{ih,t} = \text{Cov}_t\left(r_{i,t+1}, (\text{E}_{t+1} - \text{E}_t)\sum_{j=1}^{\infty}\rho^j r_{w,t+j+1}\right)$$
$$V_{i\sigma,t} = \text{Cov}_t\left(r_{i,t+1}, (\text{E}_{t+1} - \text{E}_t)\sum_{j=1}^{\infty}\rho^j \mu_{t+j+1}\right)$$

The term $V_{iw,t}$ is the covariance of asset i's return with the *current* market return. The term $V_{ih,t}$ captures the covariance with revisions in expectations of *future* market returns—a hedge demand channel that distinguishes the Epstein-Zin model from the static CAPM. The term $V_{i\sigma,t}$ measures exposure to changes in future conditional variances or risk premia.

Substituting this decomposition into (5) yields

$$\text{E}_t(r_{i,t+1}) - r_{f,t+1} = -\frac{V_{ii,t}}{2} + \gamma V_{iw,t} + (\gamma - 1)V_{ih,t} - \frac{\theta}{\psi}V_{i\sigma,t} \quad (14)$$

The first term $-V_{ii,t}/2$ is a convexity adjustment. Current market risk is priced at γ, the coefficient of relative risk aversion. News about future market returns is priced at $\gamma - 1$: an asset that covaries positively with good news about future investment opportunities provides a hedge, so it commands a lower expected return when $\gamma > 1$. The last term reflects aversion to changing second moments.

Innovations in Market Returns

An alternative decomposition, developed by Bansal and Yaron (2004), starts from the symmetric perspective of explaining *market return* innovations rather than consumption growth innovations. Rearranging (11) to solve for unexpected market returns gives

$$r_{w,t+1} - \text{E}_t\, r_{w,t+1} = \Delta c_{t+1} - \text{E}_t\, \Delta c_{t+1} + (\text{E}_{t+1} - \text{E}_t)\sum_{j=1}^{\infty}\rho^j(\Delta c_{t+j+1} - r_{w,t+j+1}) \quad (15)$$

Equation (15) shows that unexpected market returns are high when unexpected consumption growth is high, when upward revisions in expected future consumption growth are large, or when downward revisions in future market returns are large. The last source acts with the opposite sign because higher future discount rates reduce current asset prices.

Flipping (12) to express expected market returns as a function of expected consumption growth gives

$$\mathrm{E}_t\, r_{w,t+1} = \frac{1}{\psi}\left(\mathrm{E}_t\, \Delta c_{t+1} - \mu_t\right) \tag{16}$$

Substituting this into (15) to eliminate future market returns in favor of future consumption growth yields

$$\begin{aligned} r_{w,t+1} - \mathrm{E}_t\, r_{w,t+1} = \Delta c_{t+1} - \mathrm{E}_t\, \Delta c_{t+1} & \\ + \left(1 - \frac{1}{\psi}\right)(\mathrm{E}_{t+1} - \mathrm{E}_t) \sum_{j=1}^{\infty} \rho^j \Delta c_{t+j+1} & \\ + (\mathrm{E}_{t+1} - \mathrm{E}_t) \sum_{j=1}^{\infty} \rho^j \mu_{t+j+1} & \end{aligned} \tag{17}$$

Rewriting the Covariance with the Market Return. Equation (17) decomposes the conditional covariance of asset i's return with the market return as

$$\mathrm{Cov}_t(r_{i,t+1}, r_{w,t+1}) = V_{ic,t} + \left(1 - \frac{1}{\psi}\right) V_{ig,t} + V_{i\sigma,t}$$

where

$$\begin{aligned} V_{ic,t} &= \mathrm{Cov}_t(r_{i,t+1}, \Delta c_{t+1}) \\ V_{ig,t} &= \mathrm{Cov}_t\left(r_{i,t+1}, (\mathrm{E}_{t+1} - \mathrm{E}_t) \sum_{j=1}^{\infty} \rho^j \Delta c_{t+j+1}\right) \\ V_{i\sigma,t} &= \mathrm{Cov}_t\left(r_{i,t+1}, (\mathrm{E}_{t+1} - \mathrm{E}_t) \sum_{j=1}^{\infty} \rho^j \mu_{t+j+1}\right) \end{aligned}$$

Here $V_{ic,t}$ is the covariance with current consumption growth, $V_{ig,t}$ captures the covariance with revisions in expectations of *future* consumption growth—the *long-run risk* channel central to Bansal and Yaron (2004)—and $V_{i\sigma,t}$ reflects exposure to changing risk premia.

Substituting into (5) produces the risk premium in terms of consumption-based components:

$$\mathrm{E}_t(r_{i,t+1}) - r_{f,t+1} = -\frac{V_{ii,t}}{2} + \gamma V_{ic,t} + \left(\gamma - \frac{1}{\psi}\right) V_{ig,t} + (1-\theta) V_{i\sigma,t} \tag{18}$$

This is the central pricing equation of the long-run risk literature. Current consumption risk is priced at γ. Long-run consumption risk—the covariance with revisions in expected future consumption growth—is priced at $\gamma - 1/\psi$. When $\gamma > 1/\psi$, which holds whenever the agent prefers early resolution of uncertainty, assets that hedge long-run consumption risk earn lower expected returns, so long-run risk carries a positive premium. Finally, variance risk is priced at $1 - \theta$.

References

Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance* 59 (4): 1481–509.

Campbell, John Y. 1993. "Intertemporal Asset Pricing Without Consumption Data." *The American Economic Review*, 487–512.