Risk Premia in Intertemporal Models

Lorenzo Naranjo

March 2026

This notebook extends the lognormal consumption growth framework developed in Lognormal Consumption Growth to characterize risk premia in intertemporal asset pricing models, following Campbell (1993) and Bansal and Yaron (2004). The key step is to log-linearize the wealth accumulation equation, which decomposes consumption and market-return innovations into infinite discounted sums of future fundamentals. These decompositions show that risk premia extend beyond static consumption covariance: assets are also compensated for exposure to news about future market returns — a hedging demand channel emphasized by Campbell (1993) — and for exposure to long-run consumption growth risk and changing conditional variances, as in Bansal and Yaron (2004).

Throughout, lowercase letters denote logs, $\delta = -\ln\beta$, and $\theta = (1-\gamma)/(1-1/\psi)$. The lognormal framework of that notebook is maintained: consumption growth, log returns, and the log SDF are jointly conditionally normal. Two results from there serve as the point of departure. First, the lognormal pricing condition

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

holds for any SDF. Second, under Epstein-Zin preferences the log risk premium for any asset i satisfies

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

where $V_{ic,t} = \mathrm{Cov}_t(r_{i,t+1}, \Delta c_{t+1})$ and $V_{iw,t} = \mathrm{Cov}_t(r_{i,t+1}, r_{w,t+1})$.

Log-Linearization of the Budget Constraint

Campbell (1993) shows that log-linearizing the budget constraint turns the consumption-wealth ratio into a present-value relation linking current consumption to all future market returns and

consumption growth rates. This is the key step: because the consumption-wealth ratio is predetermined, any revision in expectations must leave its discounted future value unchanged, and that accounting identity is what drives the innovation decompositions throughout this section. Dividing $W_{t+1} = R_{w,t+1}(W_t - C_t)$ by W_t and taking logs gives

$$\Delta w_{t+1} = r_{w,t+1} + \ln(1 - \exp(c_t - w_t)) \tag{3}$$

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 $cw_t = c_t - w_t$.

Since the consumption-wealth ratio fluctuates around a stable long-run mean $\bar{x}$, a first-order Taylor expansion of $\ln(1 - e^x)$ around $\bar{x}$ provides a good approximation (Campbell and Shiller 1988; Campbell 1993):

$$\ln(1 - e^x) \approx \ln(1 - e^{\bar{x}}) - \frac{e^{\bar{x}}}{1 - e^{\bar{x}}}(x - \bar{x}) \tag{4}$$

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

$$\ln(1 - e^x) \approx k + \left(1 - \frac{1}{\rho}\right)x \tag{5}$$

where k can be computed from (4). The parameter ρ is slightly less than one and plays a central discounting role for future terms throughout the model. In spirit, this linearization brings the discrete-time model closer to its continuous-time counterpart, where Itô's lemma makes the budget constraint exactly linear in logs — while retaining the tractability of present-value expansions and innovation decompositions that are difficult to replicate in continuous time.

Substituting (5) into (3) 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. The appendix derives the general solution to equations of this form and establishes the revision identity used throughout — applying that result (14), we obtain

$$cw_t = \sum_{j=1}^{\infty} \rho^j (k + r_{w,t+j} - \Delta c_{t+j}) \tag{6}$$

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 (6) 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:

$$\Delta c_{t+1} - \mathsf{E}_t \Delta c_{t+1} = r_{w,t+1} - \mathsf{E}_t r_{w,t+1} + (\mathsf{E}_{t+1} - \mathsf{E}_t) \sum_{j=1}^{\infty} \rho^j (r_{w,t+j+1} - \Delta c_{t+j+1}) \tag{7}$$

Equation (7) 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.

Applying (1) to the market return, the EZ log SDF gives $m_{t+1} + r_{w,t+1} = \theta(-\delta - \Delta c_{t+1}/\psi +$

$r_{w,t+1}$), so the pricing condition reduces to

$$-\delta - \frac{1}{\psi} E_t(\Delta c_{t+1}) + E_t(r_{w,t+1}) + \frac{\theta}{2}\sigma_t^2\left(r_{w,t+1} - \frac{\Delta c_{t+1}}{\psi}\right) = 0.$$

Multiplying by ψ and solving for $E_t(\Delta c_{t+1})$ yields expected consumption growth as

$$E_t \Delta c_{t+1} = \mu_t + \psi E_t r_{w,t+1} \tag{8}$$

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 (7) to eliminate future consumption growth in favor of future market returns gives

$$\Delta c_{t+1} - E_t \Delta c_{t+1} = r_{w,t+1} - E_t r_{w,t+1}$$

$$+ (1 - \psi)(E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j r_{w,t+j+1} \tag{9}$$

$$- (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j \mu_{t+j}$$

Campbell's Pricing Equation

Equation (9) 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

$$V_{iw,t} = \text{Cov}_t(r_{i,t+1}, r_{w,t+1})$$

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

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

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 revisions in future conditional variances: because μ_{t+j} depends on $\sigma_{t+j}^2(\psi r_{w,t+j+1} - \Delta c_{t+j+1})$ via (8), revisions in μ_{t+j} are revisions in future conditional second moments, so $V_{i\sigma,t}$ is the channel through which stochastic volatility is priced.

Substituting this decomposition into (2) yields

$$\text{E}_t(r_{i,t+1}) - r_{f,t+1} = -\frac{V_{ii,t}}{2} + \gamma V_{iw,t} + (\gamma - 1)V_{ih,t} - \frac{\theta}{\psi}V_{i\sigma,t} \tag{10}$$

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

Bansal and Yaron (2004) take the symmetric perspective: rather than expressing consumption innovations in terms of return news, they express market return innovations in terms of consumption news. Their key motivation is that if consumption growth contains a small but highly persistent component — invisible in short-run data but large in present value — then long-run consumption risk can explain the equity premium without implausibly high risk aversion.

Rearranging (7) to solve for unexpected market returns gives

$$r_{w,t+1} - E_t \, r_{w,t+1} = \Delta c_{t+1} - E_t \, \Delta c_{t+1} + (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j (\Delta c_{t+j+1} - r_{w,t+j+1}) \quad (11)$$

Equation (11) 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 (8) to express expected market returns as a function of expected consumption growth gives

$$E_t \, r_{w,t+1} = \frac{1}{\psi} \, (E_t \, \Delta c_{t+1} - \mu_t)$$

Substituting this into (11) to eliminate future market returns in favor of future consumption growth yields

$$\begin{aligned}
r_{w,t+1} - E_t \, r_{w,t+1} = {}& \Delta c_{t+1} - E_t \, \Delta c_{t+1} \\
& + \left(1 - \frac{1}{\psi}\right) (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j \Delta c_{t+j+1} \\
& + \frac{1}{\psi} (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j \mu_{t+j}
\end{aligned} \quad (12)$$

Long-Run Risk Pricing

Equation (12) decomposes the conditional covariance of asset i's return with the market return as

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

where

$$V_{ic,t} = \mathrm{Cov}_t(r_{i,t+1}, \Delta c_{t+1})$$

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

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

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 revisions in future conditional variances, the same stochastic-volatility channel as in the Campbell decomposition.

Substituting into (2) produces the risk premium in terms of consumption-based components:

$$\mathrm{E}_t(r_{i,t+1}) - r_{f,t+1} = -\frac{V_{ii,t}}{2} + \gamma V_{ic,t} + \left(\gamma - \frac{1}{\psi}\right)V_{ig,t} + \frac{1-\theta}{\psi}V_{i\sigma,t} \qquad (13)$$

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 $(1-\theta)/\psi$.

Comparing the Two Decompositions

Equations (10) and (13) are dual representations of the same underlying model, not competing theories. Both start from the Epstein-Zin SDF and the same log-linearized budget constraint; they differ only in which innovation — consumption growth or market returns — is taken as the primitive. The variance risk term $V_{i\sigma,t}$ appears in both, reflecting the same source of risk priced from either angle.

The two representations connect to different empirical strategies and theoretical traditions.

Campbell (1993) organizes risk around *return* news, which links naturally to the ICAPM (Merton 1973): assets earn premia for hedging against deteriorating investment opportunities, captured by $V_{ih,t}$. This perspective motivates decomposing realized returns into cash-flow and discount-rate components, as in Campbell and Vuolteenaho (2004), who show that cash-flow beta carries a higher price of risk than discount-rate beta. Bansal and Yaron (2004), by contrast, organize risk around *consumption* news. Their key insight is that if consumption growth contains a small but persistent component, long-run consumption risk $V_{ig,t}$ is large enough to explain the equity premium without implausibly high risk aversion. Bansal et al. (2012) provide structural estimation of this channel.

A practical bridge between the two perspectives comes from empirical proxies for the consumption-wealth ratio. Lettau and Ludvigson (2001) show that the log consumption-wealth ratio — which equals cw_t up to a constant in this model — forecasts excess returns, providing a direct link between the present-value relation (6) and asset prices. An alternative empirical approach is taken by Parker and Julliard (2005), who measure consumption risk over longer horizons to capture the slow adjustment of consumption to wealth shocks, finding stronger cross-sectional pricing than short-run consumption betas deliver.

Appendix: Solving Forward Difference Equations

The log-linearization and innovation decompositions above 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

$$y_t = \sum_{j=1}^{\infty} \rho^j x_{t+j} \tag{14}$$

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:

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

Revisions of Expectations

The solution (14) 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 $\mathsf{E}_t(x_{t+j})$; 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 the main text.

References

Bansal, Ravi, Dana Kiku, and Amir Yaron. 2012. "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices." *Critical Finance Review* 1 (1): 183–221. https://doi.org/10.1561/104.00000005.

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.

Campbell, John Y., and Robert J. Shiller. 1988. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *Review of Financial Studies* 1 (3): 195–228. https://doi.org/10.1093/rfs/1.3.195.

Campbell, John Y., and Tuomo Vuolteenaho. 2004. "Bad Beta, Good Beta." *American Economic Review* 94 (5): 1249–75. https://doi.org/10.1257/0002828043052240.

Lettau, Martin, and Sydney Ludvigson. 2001. "Consumption, Aggregate Wealth, and Expected Stock Returns." *Journal of Finance* 56 (3): 815–49. https://doi.org/10.1111/0022-1082.00347.

Merton, Robert C. 1973. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 41 (5): 867–87.

Parker, Jonathan A., and Christian Julliard. 2005. "Consumption Risk and the Cross Section of Expected Returns." *Journal of Political Economy* 113 (1): 185–222. https://doi.org/10.1086/426042.