

Recursive Utility in Continuous Time

Lorenzo Naranjo

March 2026

Introduction

The [recursive preferences in a multiperiod economy](#) notebook derived the Epstein-Zin SDF in discrete time by exploiting homotheticity and Euler's theorem. The key result was

$$m_{t+1} = \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-1/\psi} \right]^\theta \left[\frac{1}{R_{t+1}^w} \right]^{1-\theta}, \quad \theta = \frac{1-\gamma}{1-1/\psi}.$$

Under power utility ($\psi = 1/\gamma$, so $\theta = 1$) the wealth return drops out entirely; under Epstein-Zin ($\theta \neq 1$) it enters as an independent pricing factor that captures the agent's concern for news about future investment opportunities.

Passing this result to continuous time requires some care. The discrete aggregator $V_t = [(1 - \beta)c_t^\rho + \beta\mu_t^\rho]^{1/\rho}$ has no direct limit when $\Delta t \rightarrow 0$ because both the discount factor $\beta = e^{-\delta\Delta t} \rightarrow 1$ and the one-period certainty-equivalent μ_t collapse in a correlated way. Duffie and Epstein (1992) resolve this by showing that Epstein-Zin preferences in continuous time are characterized by a **stochastic differential equation** for the utility process, governed by a normalized aggregator $f(c, V)$. This object plays the role of the period felicity function in additive utility, but the continuation value V_t itself enters as an argument, separating risk aversion from intertemporal substitution at every instant.

This notebook develops stochastic differential utility (SDU), derives the continuous-time Epstein-Zin aggregator and consumption first-order condition, and explains how recursive utility modifies the stochastic discount factor relative to the additive case. The continuous-time limit is both mathematically elegant and economically important: it makes precise exactly when and how γ and ψ produce independent effects on asset prices.

The general theory here stops at the recursive-utility pricing kernel. A separate notebook, [Affine Recursive Utility in Continuous Time](#), applies these results to a one-factor affine-Gaussian state model and derives the local exponential-affine approximation for the value function and the SDF.

Stochastic Differential Utility

The Aggregator

In continuous time, the utility process V_t is defined implicitly by

$$V_t = E_t \left[\int_t^\infty f(c_s, V_s) ds \right], \quad (1)$$

where the function $f(c, v)$ is the **normalized aggregator**. Equation (1) is not a closed-form definition but a functional equation: V_t must be the process such that the stochastic integral representation holds simultaneously for every t . Define $M_t = V_t + \int_0^t f(c_s, V_s) ds$. The integral representation (1) implies $M_t = E_t \left[\int_0^\infty f(c_s, V_s) ds \right]$, which is a martingale. In a Brownian filtration, the martingale representation theorem guarantees that any square-integrable martingale is a stochastic integral with respect to \mathbf{B} , so $dM_t = \sigma_t^V \cdot d\mathbf{B}_t$ for some adapted process σ_t^V . Rearranging $dM_t = dV_t + f(c_t, V_t) dt$ yields the stochastic differential equation

$$dV_t = -f(c_t, V_t) dt + \sigma_t^V \cdot d\mathbf{B}_t, \quad (2)$$

where the drift is pinned down by the aggregator and the martingale part captures unpredictable revisions to the continuation value as new information arrives.

For additive utility with felicity $u(c)$ and discount rate δ , the aggregator is

$$f(c, v) = \delta(u(c) - v).$$

To see why (1) then reduces to the standard discounted utility representation, substitute into

the functional equation to get the fixed-point condition

$$V_t = E_t \left[\int_t^\infty \delta (u(c_s) - V_s) ds \right].$$

One can verify that $V_t = E_t \left[\int_t^\infty \delta e^{-\delta(s-t)} u(c_s) ds \right]$ solves this equation: the factor δ ensures that in steady state $V = u(c)$, so the continuation value has the same units as the felicity function itself.

Substituting $f(c, v) = \delta(u(c) - v)$ into (2) gives the additive utility SDE directly:

$$dV_t = \delta(V_t - u(c_t)) dt + \sigma_t^V \cdot d\mathbf{B}_t,$$

where the martingale term $\sigma_t^V \cdot d\mathbf{B}_t$ captures revisions to V_t driven by news about the future consumption path. The steady-state condition $f = 0$ gives $V = u(c)$, consistent with the normalized integral formula.

Epstein-Zin Aggregator

To specialize the aggregator to Epstein-Zin preferences, impose two conditions on f . First, require $f(c, v) = 0$ if and only if $v = c^{1-\gamma}/(1-\gamma)$: a constant consumption path c is a steady state precisely when the continuation value equals the CRRA value of c enjoyed forever. Inverting this condition defines the certainty-equivalent consumption

$$\hat{c}(v) = [(1-\gamma)v]^{1/(1-\gamma)},$$

the unique consumption level whose CRRA lifetime utility equals v . This power transformation requires $(1-\gamma)v > 0$, so along any admissible path the utility process must have the same sign as $(1-\gamma)$: $v > 0$ when $\gamma < 1$ and $v < 0$ when $\gamma > 1$.

Second, require f to respect CRRA scale invariance. Scaling all consumption by λ scales utility by $\lambda^{1-\gamma}$, so demand $f(\lambda c, \lambda^{1-\gamma} v) = \lambda^{1-\gamma} f(c, v)$. This homogeneity condition forces f to depend on c and v only through the ratio $c/\hat{c}(v)$, with the prefactor $(1-\gamma)v$ supplying the

correct scaling:

$$f(c, v) = (1 - \gamma)v \cdot g\left(\frac{c}{\hat{c}(v)}\right),$$

for some function g with $g(1) = 0$. The ratio $c/\hat{c}(v)$ measures how current consumption compares to its break-even level, and g determines how sensitively the drift responds to deviations from that benchmark. Parametrizing g as a CES function with elasticity $\psi > 0$ gives $g(x) = \delta(x^{1-1/\psi} - 1)/(1 - 1/\psi)$ and hence the Epstein-Zin normalized aggregator

$$f(c, v) = \frac{\delta(1 - \gamma)v}{1 - 1/\psi} \left[\left(\frac{c}{\hat{c}(v)} \right)^{1-1/\psi} - 1 \right]. \quad (3)$$

The two preference parameters enter through entirely separate objects: γ determines $\hat{c}(v)$, governing how the agent values uncertainty about future utility, while ψ determines the curvature of g , governing how the drift reacts to gaps between current and break-even consumption. This separation in the aggregator is precisely what additive utility cannot achieve, even though equilibrium allocations and prices still depend jointly on the solution for the value-function coefficient.

In the limit $\psi \rightarrow 1$, the CES function $(x^{1-1/\psi} - 1)/(1 - 1/\psi)$ converges to $\ln x$, giving

$$f(c, v) \xrightarrow{\psi \rightarrow 1} \delta(1 - \gamma)v \ln\left(\frac{c}{\hat{c}(v)}\right).$$

For additive utility $\psi = 1/\gamma$, so $1 - 1/\psi = 1 - \gamma$ and $\hat{c}^{1-\gamma} = (1 - \gamma)v$. Substituting into (3) gives

$$f(c, v) = \delta \left[\frac{c^{1-\gamma}}{1 - \gamma} - v \right] = \delta(u(c) - v),$$

recovering the additive aggregator exactly. Under this normalization, the utility process V_t represents the average future felicity, which is the standard representation used to derive the Epstein-Zin SDF in continuous time.

The Bellman Equation

The agent maximizes V_0 by choosing a consumption plan $c_t \geq 0$ and a portfolio weight vector α_t . The investment opportunity set is time-varying: the short rate $r(\mathbf{z}_t)$, the vector of expected excess returns $\boldsymbol{\mu}(\mathbf{z}_t) - r(\mathbf{z}_t)\mathbf{l}$, and the return volatility matrix $\boldsymbol{\sigma}(\mathbf{z}_t)$ all depend on a state vector \mathbf{z}_t . Wealth therefore evolves as

$$dW_t = [W_t (r(\mathbf{z}_t) + \alpha_t'(\boldsymbol{\mu}(\mathbf{z}_t) - r(\mathbf{z}_t)\mathbf{l})) - c_t] dt + W_t \alpha_t' \boldsymbol{\sigma}(\mathbf{z}_t) d\mathbf{B}_t,$$

and \mathbf{z}_t follows its own diffusion driven by the same Brownian motion, capturing the covariation between portfolio returns and shifts in investment opportunities:

$$d\mathbf{z}_t = \boldsymbol{\mu}^z(\mathbf{z}_t) dt + \boldsymbol{\sigma}^z(\mathbf{z}_t) d\mathbf{B}_t.$$

Then the HJB equation for stochastic differential utility is

$$0 = \sup_{c, \alpha} \{f(c, V(W, \mathbf{z})) + \mathcal{L}^{c, \alpha} V(W, \mathbf{z})\},$$

where $\mathcal{L}^{c, \alpha}$ is the Ito generator associated with the joint process (W_t, \mathbf{z}_t) . Expanding it explicitly,

$$\begin{aligned} \mathcal{L}^{c, \alpha} V &= V_W [W(r(\mathbf{z}) + \alpha'(\boldsymbol{\mu}(\mathbf{z}) - r(\mathbf{z})\mathbf{l})) - c] + \frac{1}{2} V_{WW} W^2 \|\alpha' \boldsymbol{\sigma}(\mathbf{z})\|^2 \\ &\quad + (\nabla_z V)' \boldsymbol{\mu}^z + \frac{1}{2} \text{tr}(\boldsymbol{\sigma}^z (\boldsymbol{\sigma}^z)' H_z V) + W (\nabla_z V_W)' \boldsymbol{\sigma}^z \boldsymbol{\sigma}(\mathbf{z})' \alpha, \end{aligned}$$

where $H_z V$ is the Hessian of V with respect to \mathbf{z} and the last term captures the covariation between wealth and the state variables. Conjecturing that the value function is of the form $V(W, \mathbf{z}) = h(\mathbf{z}) W^{1-\gamma} / (1-\gamma)$ for some function $h > 0$, homotheticity implies that optimal consumption is proportional to wealth: $c/W = \kappa(\mathbf{z})$. The key derivatives are

$$V_W(W, \mathbf{z}) = h(\mathbf{z}) W^{-\gamma},$$

$$V_{WW}(W, \mathbf{z}) = -\gamma h(\mathbf{z}) W^{-\gamma-1},$$

and

$$\nabla_z V(W, \mathbf{z}) = \frac{W^{1-\gamma}}{1-\gamma} \nabla_z h(\mathbf{z}).$$

Substituting the homothetic guess into the HJB separates the choice variables from the scale variable W . The first-order condition for current consumption is

$$f_c(c, V(W, \mathbf{z})) = V_W(W, \mathbf{z}),$$

which says that the marginal gain from an extra unit of current consumption must equal the shadow value of one more unit of wealth. Using the Epstein-Zin aggregator (3) and the identity $V = hW^{1-\gamma}/(1-\gamma)$, this condition simplifies, for $\psi \neq 1$, to

$$\delta c^{-1/\psi} [(1-\gamma)V]^{1-1/\theta} = V_W, \quad \theta = \frac{1-\gamma}{1-1/\psi}.$$

To verify this, write $J = (1-\gamma)V$, so $\hat{c}(V) = J^{1/(1-\gamma)}$. Then

$$f(c, V) = \frac{\delta J}{1-1/\psi} \left[\left(\frac{c}{J^{1/(1-\gamma)}} \right)^{1-1/\psi} - 1 \right],$$

and differentiating with respect to c gives

$$f_c(c, V) = \delta c^{-1/\psi} J^{1-\frac{1-1/\psi}{1-\gamma}} = \delta c^{-1/\psi} J^{1-1/\theta}.$$

Differentiating with respect to v requires the chain rule through $\hat{c}(v)$. Since $\hat{c}(v) = [(1-\gamma)v]^{1/(1-\gamma)}$, one has $d\hat{c}/dv = \hat{c}(v)/[(1-\gamma)v]$, so $dx/dv = -x/[(1-\gamma)v]$ where $x = c/\hat{c}(v)$. Applying this to (3):

$$f_v(c, v) = \frac{\delta(1-\gamma)}{1-1/\psi} [x^{1-1/\psi} - 1] - \delta x^{1-1/\psi} = \delta \left[(\theta - 1) \left(\frac{c}{\hat{c}(v)} \right)^{1-1/\psi} - \theta \right]. \quad (4)$$

When $\theta = 1$ (power utility), $f_v = -\delta$ everywhere. When $\theta \neq 1$, f_v is state-dependent through the ratio $c/\hat{c}(v)$, which measures how far current consumption deviates from its break-even level.

After substituting $V = hW^{1-\gamma}/(1-\gamma)$ and $V_W = hW^{-\gamma}$ into the FOC and collecting powers of W , one obtains, for $\psi \neq 1$,

$$\frac{c_t}{W_t} = \delta^\psi h(\mathbf{z}_t)^{-\psi/\theta}.$$

The portfolio first-order condition (differentiating the HJB with respect to α) gives the standard Merton decomposition:

$$\alpha^* = \underbrace{\frac{1}{\gamma}(\sigma(\mathbf{z})\sigma(\mathbf{z})')^{-1}(\boldsymbol{\mu}(\mathbf{z}) - r(\mathbf{z})\mathbf{l})}_{\text{myopic demand}} + \underbrace{\frac{1}{\gamma}(\sigma(\mathbf{z})\sigma(\mathbf{z})')^{-1}\sigma(\mathbf{z})(\sigma^Z)' \frac{\nabla_{\mathbf{z}} h(\mathbf{z})}{h(\mathbf{z})}}_{\text{hedging demand}}. \quad (5)$$

The myopic component maximizes the instantaneous Sharpe ratio; the hedging component tilts the portfolio toward assets that co-vary with changes in investment opportunities. The two preference parameters enter only through $h(\mathbf{z})$: any two specifications sharing the same h yield identical portfolios. For pricing, the key point is more subtle than under additive utility. In stochastic differential utility, the intertemporal marginal rate of substitution is not simply $e^{-\delta t}V_W$. The correct marginal-utility process is

$$\Lambda_t = \Lambda_0 \exp\left(\int_0^t f_v(c_s, V_s) ds\right) f_c(c_t, V_t),$$

and using the consumption FOC this can be written equivalently as

$$\Lambda_t = \Lambda_0 \exp\left(\int_0^t f_v(c_s, V_s) ds\right) \frac{V_W(W_t, \mathbf{z}_t)}{V_W(W_0, \mathbf{z}_0)}. \quad (6)$$

Only in the additive case, where $f_v(c, v) = -\delta$, does this reduce to the familiar formula

$$\Lambda_t = \Lambda_0 e^{-\delta t} \frac{V_W(W_t, \mathbf{z}_t)}{V_W(W_0, \mathbf{z}_0)}.$$

For Epstein-Zin preferences, $f_v(c_t, V_t)$ is state dependent, so recursive utility contributes an additional discounting term through continuation-value sensitivity.

The Stochastic Discount Factor

General Form

Equation (6) is the correct starting point for asset pricing under SDU. Differentiating $\Lambda_t = \Lambda_0 e^{\int_0^t f_v ds} V_W(W_t, \mathbf{z}_t) / V_W(W_0, \mathbf{z}_0)$ gives

$$\frac{d\Lambda_t}{\Lambda_t} = f_v(c_t, V_t) dt + \frac{dV_W(W_t, \mathbf{z}_t)}{V_W(W_t, \mathbf{z}_t)}.$$

With the homothetic guess $V_W = h(\mathbf{z})W^{-\gamma}$, Ito's lemma extracts the stochastic component of dV_W/V_W :

$$\left. \frac{dV_W}{V_W} \right|_{d\mathbf{B}} = -\gamma \boldsymbol{\sigma}_t^W \cdot d\mathbf{B}_t + \frac{(\nabla_{\mathbf{z}} h(\mathbf{z}_t))' \boldsymbol{\sigma}^Z(\mathbf{z}_t)}{h(\mathbf{z}_t)} \cdot d\mathbf{B}_t,$$

where $\boldsymbol{\sigma}_t^W = \boldsymbol{\sigma}' \boldsymbol{\alpha}_t^*$ collects the wealth-return volatilities. No-arbitrage pins the drift of Λ_t to $-r_t$, so the SDF satisfies

$$\frac{d\Lambda_t}{\Lambda_t} = -r_t dt - \underbrace{\gamma \boldsymbol{\sigma}_t^W}_{\text{wealth risk}} \cdot d\mathbf{B}_t + \underbrace{\frac{(\nabla_{\mathbf{z}} h)' \boldsymbol{\sigma}^Z}{h}}_{\text{opportunity risk}} \cdot d\mathbf{B}_t. \quad (7)$$

The first diffusion term prices wealth-return shocks exactly as under power utility. The second prices news about future investment opportunities through $\nabla_{\mathbf{z}} h/h$, the logarithmic sensitivity of the value-function coefficient to state-variable shocks. The sign of this term is model-dependent: it depends on how the state vector is defined and on how shocks to \mathbf{z} affect the continuation-value coefficient h . In many applications better future opportunities raise h when $\gamma < 1$ and lower it when $\gamma > 1$ (because V is then negative), but that interpretation requires additional structure beyond the general SDU derivation.

Power Utility as a Special Case

Under power utility ($\psi = 1/\gamma$, so $\theta = 1$), the aggregator becomes additive and $f_v(c, v) = -\delta$. The general SDU pricing kernel therefore collapses to the familiar form

$$\Lambda_t = \Lambda_0 e^{-\delta t} \frac{V_W(W_t, \mathbf{z}_t)}{V_W(W_0, \mathbf{z}_0)}.$$

If, in addition, the investment opportunity set is constant so that h is constant, then $V_W = hW_t^{-\gamma}$ and the SDF satisfies

$$\frac{d\Lambda_t}{\Lambda_t} = -r dt - \gamma \boldsymbol{\sigma}^W \cdot d\mathbf{B},$$

which is exactly the power-utility SDF from the [continuous-time SDF](#) notebook.

The Epstein-Zin SDF

Under Epstein-Zin preferences with $\psi \neq 1/\gamma$ (so $\theta \neq 1$), the opportunity-risk term $(\nabla_z h)' \boldsymbol{\sigma}^z / h$ in (7) does not vanish whenever the investment opportunity set is stochastic. From (4), the state-dependent discount rate along the optimal path satisfies

$$f_v(c_t, V_t) = \delta \left[(\theta - 1) \left(\frac{c_t}{\hat{c}(V_t)} \right)^{1-1/\psi} - \theta \right],$$

which is constant only when h is constant.

Under a constant opportunity set (h constant), $\nabla_z h = 0$ and f_v is constant, so (7) reduces to $d\Lambda/\Lambda = -r dt - \gamma \boldsymbol{\sigma}^W \cdot d\mathbf{B}$ regardless of ψ . In this case EZ and power utility produce the same equity risk premium; the two parameters differ only in their implications for the risk-free rate. The full separation of γ and ψ in asset prices requires a time-varying opportunity set so that $\nabla_z h \neq 0$. That is the continuous-time channel through which recursive utility makes news about future opportunities independently priced — the exact counterpart of the wealth-return factor $1/R_{t+1}^W$ in the discrete-time SDF.

The companion notebook [Affine Recursive Utility in Continuous Time](#) implements this in a one-factor Gaussian model, where h takes an exponential-affine form and (7) yields closed-form

risk prices.

References

Duffie, Darrell, and Larry G. Epstein. 1992. "Stochastic Differential Utility." *Econometrica* 60 (2): 353–94. <https://doi.org/10.2307/2951600>.