

Optimal Portfolio and Consumption in Continuous Time

Lorenzo Naranjo

March 2026

Introduction

The [intertemporal portfolio choice](#) notebook derived optimal consumption and portfolio rules in discrete time using the Bellman equation. Continuous time offers a cleaner framework: replacing the discrete recursion with a differential equation — the Hamilton-Jacobi-Bellman (HJB) equation — allows us to characterize optimal policies in closed form and to connect the value function directly to the stochastic discount factor (SDF) and equilibrium risk premia.

This notebook solves the consumption-portfolio problem under **additive utility**: the investor maximizes the expected discounted integral of a felicity function $u(c)$. The key results are Merton's portfolio separation theorem, which decomposes the optimal portfolio into a myopic and a hedging component, and his Intertemporal CAPM (ICAPM), which links equilibrium risk premia to covariances with both aggregate wealth and state variables that shift the investment opportunity set.

The analysis proceeds in four steps. We first set up the optimization problem and write down the wealth dynamics. The Bellman principle then delivers the HJB equation, whose first-order conditions characterize optimal consumption and portfolio choice. The envelope condition establishes that the marginal value of wealth equals the marginal utility of consumption, giving the SDF. Substituting the SDF into the fundamental pricing equation yields the ICAPM.

The [next notebook](#) specializes the additive-utility problem to a one-factor Gaussian state model and solves explicitly for the value function, the wealth-consumption ratio, and the portfolio rule. The subsequent [recursive utility notebook](#) then replaces additive utility with Epstein-Zin preferences, where a modified HJB equation introduces an additional pricing factor that captures

the investor's concern for news about future investment opportunities. The structure developed here — the HJB, the envelope condition, and the market price of risk — carries over directly.

The Optimization Problem

An investor holds a portfolio of N risky assets and a money-market account. The risky assets follow the diffusion system from [Discount Factors in Continuous Time](#):

$$\frac{d\mathbf{S}}{\mathbf{S}} = \boldsymbol{\mu}(\mathbf{z}) dt + \boldsymbol{\sigma}(\mathbf{z}) d\mathbf{B},$$

where $\boldsymbol{\mu}(\mathbf{z})$ is an $N \times 1$ vector of expected returns, $\boldsymbol{\sigma}(\mathbf{z})$ is an $N \times K$ volatility matrix, and \mathbf{B} is a K -dimensional standard Brownian motion. The risk-free rate is $r(\mathbf{z})$. All coefficients depend on a vector of state variables $\mathbf{z} \in \mathbb{R}^L$ that summarizes the investment opportunity set and evolves as

$$d\mathbf{z} = \boldsymbol{\mu}^z(\mathbf{z}) dt + \boldsymbol{\sigma}^z(\mathbf{z}) d\mathbf{B}, \quad (1)$$

where $\boldsymbol{\sigma}^z(\mathbf{z})$ is an $L \times K$ matrix. The same Brownian motion \mathbf{B} drives both asset returns and state variable innovations, capturing the covariation between portfolio returns and shifts in investment opportunities.

Let $\boldsymbol{\alpha}_t$ denote the $N \times 1$ vector of portfolio weights in risky assets. The wealth process satisfies

$$\frac{dW}{W} = (r(\mathbf{z}) + \boldsymbol{\alpha}'(\boldsymbol{\mu}(\mathbf{z}) - r(\mathbf{z})\mathbf{l})) dt + \boldsymbol{\alpha}'\boldsymbol{\sigma}(\mathbf{z}) d\mathbf{B} - \frac{c}{W} dt, \quad (2)$$

where $c \geq 0$ is the consumption rate and \mathbf{l} is the $N \times 1$ vector of ones. The first term is the portfolio return net of consumption, and the second is the diffusion component carrying portfolio-return risk.

The investor maximizes expected discounted utility:

$$\max_{\{c_s, \boldsymbol{\alpha}_s\}_{s \geq t}} E_t \left[\int_t^\infty e^{-\delta(s-t)} u(c_s) ds \right], \quad (3)$$

subject to (2). We assume $u' > 0$, $u'' < 0$, and Inada conditions $u'(0) = \infty$, $u'(\infty) = 0$, which guarantee an interior optimum.

The Hamilton-Jacobi-Bellman Equation

The **value function** $V(W, \mathbf{z})$ gives the highest lifetime utility achievable from state (W, \mathbf{z}) :

$$V(W, \mathbf{z}) = \max_{\{c_s, \boldsymbol{\alpha}_s\}_{s \geq t}} E_t \left[\int_t^\infty e^{-\delta(s-t)} u(c_s) ds \right].$$

To derive its characterizing equation, apply the Bellman principle over an interval of length dt :

$$V(W, \mathbf{z}) = \max_{c, \boldsymbol{\alpha}} \{u(c) dt + e^{-\delta dt} E_t[V(W + dW, \mathbf{z} + d\mathbf{z})]\}.$$

Expanding $e^{-\delta dt} \approx 1 - \delta dt$ and applying Ito's lemma to $V(W + dW, \mathbf{z} + d\mathbf{z})$, taking expectations (so that the $d\mathbf{B}$ terms vanish), collecting terms of order dt , and rearranging yields the **Hamilton-Jacobi-Bellman (HJB) equation**:

$$\delta V = \max_{c, \boldsymbol{\alpha}} \{u(c) + \mathcal{L}^{c, \boldsymbol{\alpha}} V\}, \quad (4)$$

where $\mathcal{L}^{c, \boldsymbol{\alpha}}$ is the **Ito generator** of the joint process (W_t, \mathbf{z}_t) under controls $(c, \boldsymbol{\alpha})$:

$$\begin{aligned} \mathcal{L}^{c, \boldsymbol{\alpha}} V &= V_W [W(r + \boldsymbol{\alpha}'(\boldsymbol{\mu} - r\mathbf{1})) - c] + \frac{1}{2} V_{WW} W^2 \boldsymbol{\alpha}' \boldsymbol{\sigma} \boldsymbol{\sigma}' \boldsymbol{\alpha} \\ &\quad + (\nabla_z V)' \boldsymbol{\mu}^z + \frac{1}{2} \text{tr}(\boldsymbol{\sigma}^z (\boldsymbol{\sigma}^z)' H_z V) + W (\nabla_z V_W)' \boldsymbol{\sigma}^z \boldsymbol{\sigma}' \boldsymbol{\alpha}. \end{aligned} \quad (5)$$

Here $V_W = \partial V / \partial W$, $V_{WW} = \partial^2 V / \partial W^2$, $\nabla_z V$ is the gradient of V with respect to \mathbf{z} , $\nabla_z V_W$ is the gradient of the marginal value of wealth with respect to \mathbf{z} , and $H_z V$ is the Hessian of V with respect to \mathbf{z} . The last term in (5) captures the Ito cross-variation between wealth and state variables: since dW and $d\mathbf{z}$ share the same Brownian driver, they are instantaneously correlated.

The HJB equation (4) is the continuous-time counterpart of the discrete Bellman equation. The left-hand side δV is the required return on the value function — the rate the investor demands to postpone utility — and the right-hand side is the maximum flow of utility plus the expected capital gain on V per unit time. At the optimum, the two are equal.

First-Order Conditions

The HJB equation (4) is solved by maximizing the right-hand side over (c, α) pointwise at each state (W, \mathbf{z}) .

Consumption. Differentiating with respect to c :

$$u'(c^*) = V_W(W, \mathbf{z}). \quad (6)$$

The optimal consumption rate equates the marginal utility of consuming today to the shadow value V_W of an extra unit of wealth. This is the continuous-time envelope condition: it links preferences directly to the value function.

Portfolio. Differentiating the right-hand side of (4) with respect to α and setting equal to zero:

$$V_W W(\boldsymbol{\mu} - r\mathbf{l}) + V_{WW} W^2 \boldsymbol{\sigma} \boldsymbol{\sigma}' \boldsymbol{\alpha}^* + W \boldsymbol{\sigma} (\boldsymbol{\sigma}^z)' \nabla_z V_W = \mathbf{0}. \quad (7)$$

The three terms represent the marginal benefit of tilting toward higher-expected-return assets (first term), the variance cost of bearing more return risk (second term), and the contribution of portfolio choice to the covariance between wealth and state variable innovations (third term).

Merton's Portfolio Separation

Let $\text{rra} = -W V_{WW} / V_W > 0$ denote the coefficient of relative risk aversion implied by the value function. Substituting $V_{WW} = -\text{rra} V_W / W$ into (7) and dividing by $W V_W$ gives:

$$(\boldsymbol{\mu} - r\mathbf{l}) - \text{rra} \boldsymbol{\sigma} \boldsymbol{\sigma}' \boldsymbol{\alpha}^* + \boldsymbol{\sigma} (\boldsymbol{\sigma}^z)' \frac{\nabla_z V_W}{V_W} = \mathbf{0}.$$

Assuming $\boldsymbol{\sigma} \boldsymbol{\sigma}'$ is invertible, the optimal risky-asset portfolio is:

Property 1 (Merton's Portfolio Separation).

$$\alpha^* = \underbrace{\frac{1}{rra}(\sigma\sigma')^{-1}(\boldsymbol{\mu} - r\mathbf{1})}_{\text{myopic demand}} + \underbrace{\frac{1}{rra}(\sigma\sigma')^{-1}\boldsymbol{\sigma}(\boldsymbol{\sigma}^z)'\frac{\nabla_z V_W}{V_W}}_{\text{hedging demand}}. \quad (8)$$

The decomposition in (8) separates portfolio choice into two economically distinct motives:

1. **Myopic demand.** The first term is the portfolio that maximizes the instantaneous Sharpe ratio, scaled by the inverse of risk aversion. It is identical to the portfolio a one-period investor would choose and depends only on the current levels of $(\boldsymbol{\mu}, \boldsymbol{\sigma}, r)$, not on their dynamics.
2. **Hedging demand.** The second term arises from the investor's desire to hedge against future shifts in investment opportunities. The matrix $(\sigma\sigma')^{-1}\boldsymbol{\sigma}(\boldsymbol{\sigma}^z)'$ identifies the portfolio that best spans the Brownian innovations driving \mathbf{z} . The vector $\nabla_z V_W/V_W$ weights each state variable by the elasticity of the marginal value of wealth: when $\partial^2 V/\partial W \partial z_k > 0$, a positive shock to z_k raises V_W , so the investor overweight assets correlated with z_k to offset periods of deteriorating opportunities.

When the investment opportunity set is constant (or $\nabla_z V_W = 0$), the hedging demand vanishes and all investors hold the same risky portfolio, scaled by $1/rra$. This is the **two-fund separation** result of the static CAPM: investors choose between the risk-free asset and a single risky portfolio.

The Stochastic Discount Factor

The consumption FOC (6) establishes $u'(c^*) = V_W$. Under additive utility, the natural candidate for the SDF is therefore

$$\Lambda_t = e^{-\delta t} V_W(W_t, \mathbf{z}_t) = e^{-\delta t} u'(c_t^*), \quad (9)$$

which is the continuous-time analogue of $m_{t+1} = \beta u'(c_{t+1})/u'(c_t)$ in discrete time.

That (9) satisfies the no-arbitrage condition $E(d\Lambda/\Lambda) = -r dt$ follows from applying Ito's lemma to V_W and substituting the HJB equation into the resulting drift, which equals $-r$ at the

optimum. The diffusion component of $d\Lambda/\Lambda$ determines the market price of risk. Applying Ito's lemma to $\Lambda = e^{-\delta t}V_W$, the $d\mathbf{B}$ part is:

$$\frac{d\Lambda}{\Lambda} \Big|_{d\mathbf{B}} = \frac{V_{WW}}{V_W} dW \Big|_{d\mathbf{B}} + \frac{(\nabla_z V_W)'}{V_W} dz \Big|_{d\mathbf{B}} = \left(-rra \boldsymbol{\alpha}^{*'} \boldsymbol{\sigma} + \frac{(\nabla_z V_W)'}{V_W} \boldsymbol{\sigma}^z \right) d\mathbf{B}.$$

Comparing with the generic SDF form $d\Lambda/\Lambda = -r dt - \boldsymbol{\lambda}' d\mathbf{B}$ from [Discount Factors in Continuous Time](#), the market price of risk vector is

$$\boldsymbol{\lambda}' = rra \boldsymbol{\alpha}^{*'} \boldsymbol{\sigma} - \frac{(\nabla_z V_W)'}{V_W} \boldsymbol{\sigma}^z. \quad (10)$$

The first term is the compensation per unit of exposure to each Brownian shock, driven by the portfolio's volatility weighted by risk aversion. The second term is the compensation for bearing state-variable risk.

The Intertemporal CAPM

The no-arbitrage condition requires that for any risky asset S paying a dividend yield D/S ,

$$\mathbb{E} \left(\frac{dS}{S} \right) + \frac{D}{S} dt - r dt = -\frac{d\Lambda}{\Lambda} \frac{dS}{S}. \quad (11)$$

Substituting the SDF dynamics from (10) into the right-hand side, and using $(d\mathbf{B})(d\mathbf{B})' = \mathbf{I} dt$ to express $\boldsymbol{\alpha}^{*'} \boldsymbol{\sigma} d\mathbf{B}$ and $\boldsymbol{\sigma}^z d\mathbf{B}$ as the diffusion parts of dW/W and dz , yields Merton's Intertemporal CAPM.

Property 2 (Merton's Intertemporal CAPM). *The equilibrium risk premium of any risky asset satisfies*

$$\mathbb{E} \left(\frac{dS}{S} \right) + \frac{D}{S} dt - r dt = rra \cdot \frac{dW}{W} \frac{dS}{S} - \frac{V_{Wz'}(W, \mathbf{z})}{V_W(W, \mathbf{z})} \left(dz \frac{dS}{S} \right). \quad (12)$$

Expected excess returns depend on $K + 1$ factors: the wealth portfolio and K state-variable hedging portfolios, one for each source of time-variation in the investment opportunity set.

The risk premium has two economically distinct components:

1. **Wealth risk premium.** The term $\text{rra} \cdot (dW/W)(dS/S)/dt$ is the instantaneous covariance of asset returns with wealth growth, scaled by the coefficient of relative risk aversion. Assets that covary positively with aggregate wealth are risky and command higher expected returns.
2. **Hedging demand premium.** The term $-(V_{Wz'}/V_W)(dz)(dS/S)/dt$ reflects investors' desire to hedge against shifts in the investment opportunity set. An asset that covaries positively with a state variable z_k for which $V_{Wz_k} > 0$ acts as a hedge against deteriorating investment opportunities and therefore commands a lower risk premium.

When there are no state variables — or whenever $V_{Wz_k} = 0$ for all k — the ICAPM reduces to a single-factor model in which the market (wealth) portfolio is the only priced risk. In that special case, equation (12) recovers the continuous-time CAPM: risk premiums are proportional to covariance with aggregate wealth.

The connection to the next two notebooks is direct. The [one-factor lognormal sequel](#) specializes the general HJB to a Gaussian state variable and solves exactly for the value function, the consumption-wealth ratio, and the portfolio rule. The [recursive utility notebook](#) then replaces additive utility with an Epstein-Zin aggregator in which the discount rate depends on the ratio of current consumption to its certainty-equivalent level, introducing an additional pricing factor for news about future investment opportunities.

Power Utility

Having established that $\Lambda_t = e^{-\delta t} u'(c_t^*)$ from (9), we can apply Ito's lemma directly to find the dynamics of Λ in terms of consumption growth. For a general felicity function,

$$\begin{aligned} d\Lambda &= \frac{\partial \Lambda}{\partial c} dc + \frac{1}{2} \frac{\partial^2 \Lambda}{\partial c^2} (dc)^2 + \frac{\partial \Lambda}{\partial t} dt \\ &= e^{-\delta t} u''(c) dc + \frac{1}{2} e^{-\delta t} u'''(c) (dc)^2 - \delta e^{-\delta t} u'(c) dt, \end{aligned}$$

or

$$\frac{d\Lambda}{\Lambda} = -\delta dt + \frac{1}{2} \frac{c^2 u'''(c)}{u'(c)} \left(\frac{dc}{c} \right)^2 + \frac{c u''(c)}{u'(c)} \frac{dc}{c}.$$

For power utility we have that

$$\frac{d\Lambda}{\Lambda} = -\delta dt + \frac{1}{2}\gamma(\gamma + 1)\left(\frac{dc}{c}\right)^2 - \gamma\frac{dc}{c}.$$

Write consumption growth as

$$\frac{dc}{c} = \mu_c dt + \sigma_c dB_c.$$

Substituting yields

$$\frac{d\Lambda}{\Lambda} = \left(-\delta + \frac{1}{2}\gamma(\gamma + 1)\sigma_c^2 - \gamma\mu_c\right)dt - \gamma\sigma_c dB_c. \quad (13)$$

The instantaneous risk-free rate equals minus the drift of $d\Lambda/\Lambda$:

$$r = \delta + \gamma\mu_c - \frac{1}{2}\gamma(\gamma + 1)\sigma_c^2.$$

The expression has three components. Real interest rates are high when impatience (δ) is high, since more impatient investors demand a high return to save. They are high when expected consumption growth (μ_c) is high, since agents expecting rising consumption need to save less, pushing bond prices down. Finally, they are low when consumption volatility (σ_c) is high — the **precautionary savings** effect: more uncertain future consumption raises the demand for safe assets, pushing bond prices up and yields down.

We can use this to understand the **equity premium puzzle**. Consider an asset paying a dividend flow $D dt$ and following a diffusion

$$\frac{dS}{S} = \mu_S dt + \sigma_S dB_S$$

such that $(dB_S)(dB_C) = \rho dt$. Equations (11) and (13) imply

$$\mu_S + D/S - r = \gamma\rho\sigma_c\sigma_S.$$

In this simple model with power utility, the risk premium of any risky asset is higher when risk aversion (γ) is high and/or the covariance of asset returns and consumption growth is high.

Since $|\rho| \leq 1$, the previous expression implies

$$\left| \frac{\mu_S + D/S - r}{\sigma_S} \right| \leq \gamma \sigma_c.$$

In the data, the Sharpe ratio of the market is around 0.5 whereas the standard deviation of consumption growth is around 0.01. We need a RRA coefficient of at least 50 to explain the risk premium of the market. To resolve this paradox, researchers have introduced preferences that generate a more volatile stochastic discount factor, such as recursive Epstein-Zin preferences or habits.