

Intertemporal Portfolio Choice

Lorenzo Naranjo

March 2026

Introduction

The two-date portfolio problem studied in the previous notebook is a useful starting point, but it misses an important feature of real investment decisions: the investment opportunity set changes over time. Interest rates fluctuate, risk premia vary, and volatility clusters. A forward-looking investor cares not only about the return on her portfolio today, but also about how the portfolio performs when future investment opportunities are favorable or unfavorable.

This notebook extends the portfolio choice problem to an infinite horizon, where asset returns depend on a vector of **state variables** z_t that evolves over time. The state variables summarize everything relevant about the current investment opportunity set — dividend yields, yield spreads, volatility indices, or any other variable with predictive content for future returns.

The extension has two major consequences. First, the stochastic discount factor (SDF) now depends on both the level of wealth and the state of the investment opportunity set. Second, the optimal portfolio is no longer a pure mean-variance portfolio: agents also take positions to **hedge** against adverse shifts in the investment opportunity set, a phenomenon known as the **hedging demand** component of Merton's ICAPM (Merton 1973).

We proceed in four steps. We begin by deriving the Bellman equation for the infinite-horizon problem and its first-order conditions. The envelope condition then establishes that the marginal value of wealth equals marginal utility, delivering the SDF. A linear approximation of the value function yields a tractable SDF, a multi-factor risk premium formula, and an explicit expression for the optimal portfolio.

The Infinite-Horizon Problem

An investor with wealth W_t chooses a consumption plan $\{c_{t+i}\}_{i=0}^{\infty}$ and a portfolio plan $\{\alpha_{t+i}\}_{i=0}^{\infty}$ to maximize expected discounted utility:

$$\max_{\{c_{t+i}, \alpha_{t+i}\}_{i=0}^{\infty}} E_t \left(\sum_{i=0}^{\infty} \beta^i u(c_{t+i}) \right),$$

subject to the wealth accumulation equation

$$W_{t+1} = r_{t+1}^w (W_t - c_t), \quad r_{t+1}^w = \alpha_t' (\mathbf{r}_{t+1} - r_{t+1}^f \mathbf{l}) + r_{t+1}^f,$$

Here α_t is the vector of risky-asset portfolio weights, $\mathbf{r}_{t+1}$ is the vector of risky gross returns, and r_{t+1}^f is the gross risk-free return. Current savings $W_t - c_t$ are invested between t and $t + 1$, so next period's wealth is the only resource available for future consumption.

The investment opportunity set varies over time through a vector of **state variables** z_t . These state variables evolve exogenously and affect the conditional distribution of returns. For example, risky returns may satisfy

$$\mathbf{r}_{t+1} = E(\mathbf{r}_{t+1} | z_t) + \sigma(\mathbf{r}_{t+1} | z_t) \varepsilon_{t+1},$$

where ε_{t+1} is a mean-zero innovation. The relevant state at date t is therefore the pair (W_t, z_t) : wealth summarizes current resources, while z_t summarizes current investment opportunities.

Bellman Equation

The **value function** gives the highest lifetime utility attainable from a given state (W_t, z_t) :

$$V(W_t, z_t) = \max_{\{c_{t+i}, \alpha_{t+i}\}_{i=0}^{\infty}} E_t \left(\sum_{i=0}^{\infty} \beta^i u(c_{t+i}) \right),$$

where the maximum is over all feasible consumption and portfolio policies from date t onward.

To obtain the recursive formulation, separate current utility from continuation utility:

$$V(W_t, z_t) = \max_{\{c_t, \alpha_t\}} \left\{ u(c_t) + \max_{\{c_{t+1+i}, \alpha_{t+1+i}\}_{i=0}^{\infty}} E_t \left(\sum_{i=1}^{\infty} \beta^i u(c_{t+i}) \right) \right\}.$$

Re-index the continuation sum and apply the **law of iterated expectations**, $E_t[\cdot] = E_t[E_{t+1}[\cdot]]$:

$$\max_{\{c_{t+1+i}, \alpha_{t+1+i}\}_{i=0}^{\infty}} E_t \left(\sum_{i=1}^{\infty} \beta^i u(c_{t+i}) \right) = E_t \left[\beta \max_{\{c_{t+1+i}, \alpha_{t+1+i}\}_{i=0}^{\infty}} E_{t+1} \left(\sum_{i=0}^{\infty} \beta^i u(c_{t+1+i}) \right) \right].$$

After c_t and α_t are chosen, the problem at $t + 1$ has the same form as the original one, but starts from the random state (W_{t+1}, z_{t+1}) . The inner maximization is therefore $V(W_{t+1}, z_{t+1})$, and the **Bellman equation** becomes

$$V(W_t, z_t) = \max_{\{c_t, \alpha_t\}} \{ u(c_t) + \beta E_t [V(W_{t+1}, z_{t+1})] \},$$

subject to

$$W_{t+1} = r_{t+1}^w (W_t - c_t), \quad r_{t+1}^w = \alpha_t' (r_{t+1} - r_{t+1}^f \mathbf{l}) + r_{t+1}^f.$$

The Bellman equation compresses the infinite-horizon problem into a sequence of identical one-period decisions. At each date, the investor chooses current consumption and portfolio weights, recognizing that today's choices determine tomorrow's state.

First-Order Conditions

The Bellman equation has two control variables, current consumption c_t and portfolio weights α_t . Differentiating with respect to each control, and using $\partial W_{t+1} / \partial c_t = -r_{t+1}^w$ and $\partial W_{t+1} / \partial \alpha_t = (r_{t+1} - r_{t+1}^f \mathbf{l})(W_t - c_t)$, gives the first-order conditions:

FOC for consumption:

$$u'(c_t) = E_t [\beta r_{t+1}^w V_W(W_{t+1}, z_{t+1}) \mid z_t],$$

FOC for portfolio weights:

$$E_t \left[\beta (r_{t+1} - r_{t+1}^f \mathbf{1}) (W_t - c_t) V_W(W_{t+1}, z_{t+1}) \mid z_t \right] = \mathbf{0},$$

Here $V_W \equiv \partial V / \partial W$ denotes the marginal value of wealth. In an interior solution, savings $W_t - c_t$ are positive, so this common factor can be divided out of the portfolio condition:

$$E_t \left[\beta r_{t+1}^e V_W(W_{t+1}, z_{t+1}) \mid z_t \right] = \mathbf{0},$$

where $r_{t+1}^e = r_{t+1} - r_{t+1}^f \mathbf{1}$ is the vector of excess returns.

These two conditions have a simple interpretation. The consumption FOC says that the marginal utility lost by consuming one less unit today must equal the expected discounted marginal value of the extra wealth carried into tomorrow. The portfolio FOC says that, at the optimum, no marginal reallocation toward any risky asset can improve lifetime utility once returns are weighted by the future marginal value of wealth.

Envelope Condition

Let $(c_t^*(W_t, z_t), \alpha_t^*(W_t, z_t))$ denote the optimal policies. Substituting them into the Bellman equation gives

$$V(W_t, z_t) = u(c_t^*) + E_t \left[\beta V(r_{t+1}^{W*}(W_t - c_t^*), z_{t+1}) \right].$$

To see how the value function responds to an extra dollar of wealth, differentiate both sides with

respect to W_t :

$$\begin{aligned}
V_W(W_t, z_t) &= u'(c_t^*) \frac{\partial c_t^*}{\partial W_t} + E_t \left[\beta V_W(W_{t+1}, z_{t+1}) \left(\frac{\partial \alpha_t^{*'}}{\partial W_t} r_{t+1}^e (W_t - c_t^*) + r_{t+1}^{w*} \left(1 - \frac{\partial c_t^*}{\partial W_t} \right) \right) \middle| z_t \right] \\
&= \underbrace{\left(u'(c_t^*) - E_t [\beta r_{t+1}^{w*} V_W(W_{t+1}, z_{t+1})] \right)}_{= 0 \text{ by FOC for } c_t} \frac{\partial c_t^*}{\partial W_t} \\
&\quad + \beta \frac{\partial \alpha_t^{*'}}{\partial W_t} (W_t - c_t^*) \underbrace{E_t [V_W(W_{t+1}, z_{t+1}) r_{t+1}^e]}_{= 0 \text{ by FOC for } \alpha_t} + E_t [\beta V_W(W_{t+1}, z_{t+1}) r_{t+1}^{w*}].
\end{aligned}$$

The first two terms vanish by the first-order conditions, leaving the **envelope condition**:

$$V_W(W_t, z_t) = E_t [\beta V_W(W_{t+1}, z_{t+1}) r_{t+1}^{w*}] = u'(c_t^*(W_t, z_t)).$$

This is the same logic as in the two-date problem. At the optimum, the marginal value of one more dollar of wealth equals the marginal utility of consuming that dollar today. The middle expression shows the same dollar viewed as an investment: its value is the expected discounted marginal utility tomorrow, scaled by the portfolio return.

Stochastic Discount Factor

The envelope condition suggests the natural one-period pricing kernel:

$$m_{t+1} = \beta \frac{V_W(W_{t+1}, z_{t+1})}{V_W(W_t, z_t)} = \beta \frac{u'(c_{t+1}^*)}{u'(c_t^*)}.$$

The first-order conditions imply that m_{t+1} prices excess returns. From the portfolio FOC,

$$E_t [m_{t+1} r_{t+1}^e | z_t] = 0 \implies E_t [m_{t+1} r_{t+1} | z_t] = r_{t+1}^f E_t [m_{t+1} | z_t] \mathbf{1}.$$

The consumption FOC and the envelope condition also imply

$$E_t [m_{t+1} r_{t+1}^w | z_t] = 1.$$

Since the risk-free asset is also traded, it must satisfy the same pricing equation:

$$E_t [m_{t+1} r_{t+1}^f | z_t] = r_{t+1}^f E_t [m_{t+1} | z_t] = 1.$$

Therefore

$$E_t [m_{t+1} r_{t+1} | z_t] = 1,$$

so m_{t+1} is a valid SDF. The representation $m_{t+1} = \beta u'(c_{t+1})/u'(c_t)$ is the familiar consumption-based pricing kernel of the intertemporal asset-pricing literature (Lucas 1978; Breeden 1979). The equivalent form $m_{t+1} = \beta V_W(W_{t+1}, z_{t+1})/V_W(W_t, z_t)$ is often more convenient in dynamic settings because it expresses prices directly in terms of the value function.

Linearization

To move beyond the formal Bellman equation, we need an approximation for the marginal value of wealth. Rather than solve for V in closed form, we linearize V_W around the current state (W_t, z_t) :

$$V_W(W_{t+1}, z_{t+1}) \approx V_W(W_t, z_t) + V_{WW}(W_t, z_t)(W_{t+1} - W_t) + V_{Wz}(W_t, z_t)'(z_{t+1} - z_t),$$

where $V_{WW} = \partial^2 V / \partial W^2$ and $V_{Wz} = \partial^2 V / \partial W \partial z$ is the vector of cross-derivatives with respect to wealth and the state variables.

Divide through by $V_W(W_t, z_t)$ and introduce three objects that summarize the local behavior of the value function:

- **Relative risk aversion:** $rra_t = -\frac{W_t V_{WW}(W_t, z_t)}{V_W(W_t, z_t)} \geq 0$, which captures the curvature of indirect utility.
- **Consumption-to-wealth ratio:** $\delta_t = c_t/W_t \in (0, 1)$.
- **Hedging demand coefficients:** $h_t = -V_{Wz}(W_t, z_t)/V_W(W_t, z_t)$, which measure how the marginal value of wealth responds to movements in the state variables.

The wealth dynamics translate the Taylor expansion into return space. Since $W_{t+1} = r_{t+1}^w (W_t - c_t) = r_{t+1}^w W_t (1 - \delta_t)$,

$$\frac{W_{t+1} - W_t}{W_t} = r_{t+1}^w (1 - \delta_t) - 1.$$

Substituting into the definition of the SDF yields the **linear SDF approximation**

$$m_{t+1} \approx \beta \left[1 - rra_t (r_{t+1}^w (1 - \delta_t) - 1) - h_t' \Delta z_{t+1} \right],$$

where $\Delta z_{t+1} = z_{t+1} - z_t$. This approximation makes the economics transparent. The SDF moves with two kinds of shocks: shocks to next-period wealth, scaled by risk aversion, and shocks to the state variables, scaled by the hedging coefficients. If the opportunity set were constant so that z_t did not move, the second channel would disappear and the model would collapse to a single-factor specification.

This representation immediately delivers a covariance decomposition for any asset return r_{t+1}^i :

$$\text{Cov}_t(m_{t+1}, r_{t+1}^i) = \beta \left[-rra_t (1 - \delta_t) \text{Cov}_t(r_{t+1}^w, r_{t+1}^i) - h_t' \text{Cov}_t(\Delta z_{t+1}, r_{t+1}^i) \right].$$

Risk Premia and Portfolio Demand

Multi-Factor Risk Premia

The SDF pricing condition $E_t[m_{t+1} r_{t+1}^{e,i}] = 0$ and the covariance decomposition above imply

$$E_t r_{t+1}^{e,i} = \beta_t^{w,i} \lambda_t^w + \left(\beta_t^{z,i} \right)' \lambda_t^h,$$

where the **betas** measure asset i 's exposure to each source of risk,

$$\beta_t^{w,i} = \frac{\text{Cov}_t(r_{t+1}^i, r_{t+1}^w)}{\sigma_t^2(r_{t+1}^w)}, \quad \beta_t^{z,i} = \frac{\text{Cov}_t(r_{t+1}^i, \Delta z_{t+1})}{\sigma_t^2(\Delta z_{t+1})},$$

and the **risk prices** reward each unit of exposure:

$$\lambda_t^w = \beta r r a_t (1 - \delta_t) \sigma_t^2 (r_{t+1}^w), \quad \lambda_t^h = \beta h_t \sigma_t^2 (\Delta z_{t+1}).$$

This is Merton's **Intertemporal CAPM (ICAPM)** (Merton 1973): expected excess returns are driven by two types of betas. The first, $\beta_t^{w,i}$, measures the traditional covariance with the market portfolio (the “wealth beta”). The second, $\beta_t^{z,i}$, measures covariance with state variable innovations — assets that hedge against deteriorating investment opportunities earn lower expected returns because investors value them precisely when opportunities are scarce. In discrete-time empirical work, this state-variable perspective is closely related to the intertemporal pricing formulation in Campbell (1993).

Optimal Portfolio Demand

To find the optimal portfolio α_t , note that $\text{Cov}_t(r_{t+1}^w, r_{t+1}^i) = \Sigma_t \alpha_t$ where $\Sigma_t = \text{Cov}_t(\mathbf{r}_{t+1})$ is the conditional covariance matrix of risky returns. Let $\mu_t^e = E_t(r_{t+1}^e)$ and $\Sigma_t^{z,r} = \text{Cov}_t(\Delta z_{t+1}, \mathbf{r}_{t+1})$ (a matrix whose rows correspond to state variables and columns to assets). The expected excess return can be written in vector form as

$$\mu_t^e = \beta r r a_t (1 - \delta_t) \Sigma_t \alpha_t + \beta \Sigma_t^{z,r'} h_t.$$

Solving for the portfolio weights gives

$$\alpha_t = \frac{1}{\beta r r a_t (1 - \delta_t)} (\Sigma_t^{-1} \mu_t^e - \beta \Sigma_t^{-1} \Sigma_t^{z,r'} h_t).$$

This expression separates naturally into a mean-variance component and a hedging component. The first term,

$$\frac{1}{\beta r r a_t (1 - \delta_t)} \Sigma_t^{-1} \mu_t^e,$$

is the familiar tangency-portfolio demand, scaled inversely by risk aversion and by the share of wealth that is saved. It is exactly the term that would survive in a static model with a fixed investment opportunity set.

The second term,

$$-\frac{1}{rra_t(1 - \delta_t)} \Sigma_t^{-1} \Sigma_t^{z,r'} h_t,$$

is the hedging demand. It reflects positions taken not because they have high expected returns, but because they help offset adverse movements in the state variables. The vector $\Sigma_t^{-1} \Sigma_t^{z,r'} h_t$ identifies the portfolio that best spans innovations in investment opportunities, while the sign and magnitude of h_t determine how aggressively the investor wants that hedge. If $h_t = \mathbf{0}$ — either because there are no state variables or because $V_{WZ} = 0$ — this term disappears and the standard mean-variance solution is recovered.

References

- Breeden, Douglas T. 1979. “An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities.” *Journal of Financial Economics* 7 (3): 265–96.
- Campbell, John Y. 1993. “Intertemporal Asset Pricing Without Consumption Data.” *American Economic Review*, 487–512.
- Lucas, Robert E. 1978. “Asset Prices in an Exchange Economy.” *Econometrica* 46 (6): 1429–45.
- Merton, Robert C. 1973. “An Intertemporal Capital Asset Pricing Model.” *Econometrica* 41 (5): 867–87.