

# Portfolio Choice

Lorenzo Naranjo

March 2026

## Introduction

A central question in asset pricing is: what determines the price of a risky asset? One approach is to ask what a rational, utility-maximizing agent would be willing to pay. If we can characterize her optimal portfolio choice, we can read off the pricing rule directly from the first-order conditions.

This notebook takes that approach using **dynamic programming**, a framework for solving sequential optimization problems. Dynamic programming is built around three ingredients:

- **State variable.** A variable  $s$  that summarizes everything relevant about the agent's situation at the time of a decision. The value function  $V(s)$  records the maximum expected utility the agent can achieve from state  $s$  onward.
- **Control variables.** The agent's decision variables — the quantities she chooses to maximize  $V(s)$ .
- **Bellman equation.** The recursive relationship  $V(s) = \max_{\text{controls}} \{\text{current payoff} + \text{continuation value}\}$ , which breaks a multi-period problem into a sequence of simpler one-period problems.

The power of dynamic programming comes from the **envelope theorem**: differentiating  $V$  with respect to the state at the optimum eliminates all terms involving how controls respond to the state, leaving  $V'(s)$  equal to the direct effect of the state on the objective. The economic intuition is that at the optimum the agent is indifferent at the margin between consuming a dollar today and investing it — if one use were strictly better, she would have reallocated already. This indifference means the marginal value of wealth  $V'(W_0)$  equals the marginal utility of consumption  $u'(c_0)$ , the key building block of the SDF.

We proceed in two steps. We begin with the simplest dynamic programming problem — portfolio choice with no current consumption — where wealth  $W_0$  is the only state and the portfolio  $\alpha$  is the only control. The Bellman equation collapses to  $V(W_0) = \max_{\alpha} E[u(r^w W_0)]$ , and the FOC immediately delivers an SDF. We then add a second control — current consumption  $c_0$  — obtaining the canonical two-date dynamic programming problem. Its envelope condition,  $V'(W_0) = u'(c_0)$ , links the derivative of the value function to marginal utility and is the key building block for extending to longer horizons.

## Portfolio Choice with No Current Consumption

Consider an agent with initial wealth  $W_0$  who invests everything in a portfolio  $\alpha$  and consumes the proceeds at the end of the period. The **state variable** is  $W_0$ , the **control variable** is the portfolio  $\alpha$ , and the **value function**  $V(W_0)$  is the maximum expected utility attainable from state  $W_0$ .

Let  $\alpha$  denote the vector of portfolio weights in the risky assets, so that  $1 - \alpha' \mathbf{1}$  is the weight in the risk-free asset. The portfolio return is

$$r^w = \alpha' \mathbf{r} + (1 - \alpha' \mathbf{1}) r^f = \alpha' (\mathbf{r} - r^f \mathbf{1}) + r^f = \alpha' \mathbf{r}^e + r^f,$$

where  $\mathbf{r}^e = \mathbf{r} - r^f \mathbf{1}$  is the vector of excess returns. Since end-of-period wealth and consumption satisfy  $W = r^w W_0$  and  $c = W$ , the Bellman equation collapses to a single-period problem:

$$V(W_0) = \max_{\alpha} E(u(r^w W_0)).$$

The first-order condition with respect to  $\alpha$  is

$$E(u'(W) \mathbf{r}^e) = \mathbf{0}.$$

This states that, at the optimum, the expected marginal utility of wealth times each excess return is zero. Expanding:

$$E(u'(W) \mathbf{r}) = E(u'(W)) r^f \mathbf{1},$$

which can be rewritten as

$$E\left(\frac{u'(W)}{r^f E(u'(W))} \mathbf{r}\right) = \mathbf{1}.$$

A **stochastic discount factor** (SDF) is a random variable  $m$  that prices all assets:  $E(m \mathbf{r}) = \mathbf{1}$ , meaning the expected discounted return on every asset equals one. The previous expression shows that

$$m = \frac{u'(W)}{r^f E(u'(W))}$$

is a valid SDF. Indeed,  $E(m \mathbf{r}) = \mathbf{1}$  follows directly from the FOC, and  $E(m r^f) = 1$  holds since  $r^f$  is risk-free.

Define the indirect utility (value) function  $V(W_0) = E(u(r^w W_0))$ . Assuming that the optimal portfolio  $\boldsymbol{\alpha}^*(W_0)$  is differentiable in  $W_0$ , differentiating  $V$  gives

$$\begin{aligned} V'(W_0) &= E(u'(W) r^w) + W_0 \frac{d\boldsymbol{\alpha}^*}{dW_0} \underbrace{E(u'(W) \mathbf{r}^e)}_{= 0 \text{ by FOC}} \\ &= E(u'(W) r^w) \\ &= r^f E(u'(W)), \end{aligned}$$

where the last equality uses the FOC. Therefore, the SDF can equivalently be written as

$$m = \frac{u'(W)}{V'(W_0)}.$$

**Example 1** (CARA Utility). Combining CARA utility with normally distributed returns is a classical approach to portfolio choice and is one way to derive the CAPM, as shown in the Portfolio Frontier Mathematics notebooks. Let  $u(c) = -e^{-\gamma c}$  with  $\gamma > 0$ . For CARA utility it is more natural to work with dollar amounts rather than portfolio weights, so let  $y^d$  and  $\mathbf{y}$  denote the dollar amounts held in the risk-free asset and in risky assets, respectively, with  $W_0 = y^d + \mathbf{y}'\mathbf{1}$  and

$$c = W = y^d r^f + \mathbf{y}'\mathbf{r}.$$

Assuming  $\mathbf{r} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ , we have  $W \sim \mathcal{N}(y^d r^f + \mathbf{y}'\boldsymbol{\mu}, \mathbf{y}'\boldsymbol{\Sigma}\mathbf{y})$ . The expected utility is

$$\begin{aligned} E(u(W)) &= -\exp\left(-\gamma(y^d r^f + \mathbf{y}'\boldsymbol{\mu}) + \frac{\gamma^2}{2}\mathbf{y}'\boldsymbol{\Sigma}\mathbf{y}\right) \\ &= -\exp\left(-\gamma((W_0 - \mathbf{y}'\mathbf{1})r^f + \mathbf{y}'\boldsymbol{\mu}) + \frac{\gamma^2}{2}\mathbf{y}'\boldsymbol{\Sigma}\mathbf{y}\right) \\ &= -\exp\left(-\gamma(W_0 r^f + \mathbf{y}'\boldsymbol{\mu}^e) + \frac{\gamma^2}{2}\mathbf{y}'\boldsymbol{\Sigma}\mathbf{y}\right). \end{aligned}$$

The FOC with respect to  $\mathbf{y}$  gives

$$-\gamma\boldsymbol{\mu}^e + \gamma^2\boldsymbol{\Sigma}\mathbf{y} = \mathbf{0} \Rightarrow \mathbf{y} = \frac{1}{\gamma}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}^e.$$

This is the **mean-variance optimal portfolio**: the agent invests a dollar amount proportional to  $\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}^e$  in each risky asset. Moreover, the excess return vector satisfies

$$\boldsymbol{\mu}^e = \gamma\boldsymbol{\Sigma}\mathbf{y} = \gamma \text{Cov}(\mathbf{r}, \mathbf{r}'\mathbf{y}) = \gamma \text{Cov}(\mathbf{r}, r^w W_0) = \gamma W_0 \text{Cov}(\mathbf{r}, r^w),$$

so the risk premium on each asset is proportional to its covariance with the portfolio return.  $\square$

## Portfolio Optimization with Consumption and Savings

We now add a layer of complexity. The agent has two dates and must decide both how much to consume today and how to invest the remainder. The **state variable** is still initial wealth  $W_0$ , but there are now two **control variables**: current consumption  $c_0$  and portfolio weights  $\boldsymbol{\alpha}_0$ . Savings  $W_0 - c_0$  are invested and generate next-period wealth  $W_1 = r^w(W_0 - c_0)$ , which is consumed in full. The **Bellman equation** is

$$V(W_0) = \max_{\{c_0, \boldsymbol{\alpha}_0\}} u(c_0) + \beta E(u(c_1))$$

subject to

$$W_1 = r^w(W_0 - c_0), \quad c_1 = W_1, \quad r^w = \boldsymbol{\alpha}'_0(\mathbf{r} - r^f \mathbf{1}) + r^f.$$

## First-Order Conditions

The FOC with respect to  $c_0$  and  $\alpha_0$  are, respectively,

$$u'(c_0) + \beta E(u'(c_1) \cdot (-r^w)) = 0,$$

$$\beta E(u'(c_1) (\mathbf{r} - r^f \mathbf{1})) = \mathbf{0}.$$

Combining the two conditions we can rearrange to obtain

$$E\left(\beta \frac{u'(c_1)}{u'(c_0)} \mathbf{r}\right) = \mathbf{1},$$

which shows that

$$m = \beta \frac{u'(c_1)}{u'(c_0)} \in \mathcal{M}$$

is a valid SDF.

## Envelope Condition

How does the value function respond to a small increase in initial wealth? Naively, differentiating  $V(W_0)$  with respect to  $W_0$  generates terms from how the optimal choices  $c_0^*$  and  $\alpha_0^*$  adjust. But at the optimum, the agent has already balanced consumption against investment: a marginal reallocation between the two has no first-order effect on utility. Those terms therefore vanish, and only the direct effect of wealth on the objective survives.

Formally, substituting the optimal policies into the Bellman equation gives

$$V(W_0) = u(c_0^*(W_0)) + \beta E(u((\alpha_0^*(W_0)' \mathbf{r}^e + r^f)(W_0 - c_0^*(W_0)))).$$

Differentiating with respect to  $W_0$  and using the chain rule:

$$\begin{aligned} V'(W_0) &= u'(c_0) \frac{dc_0^*}{dW_0} + \beta \mathbb{E} \left[ u'(c_1) \left( r^w \left( 1 - \frac{dc_0^*}{dW_0} \right) + (W_0 - c_0) \frac{d\alpha_0^*}{dW_0} r^e \right) \right] \\ &= \underbrace{(u'(c_0) - \beta \mathbb{E}[u'(c_1) r^w])}_{= 0 \text{ by FOC for } c_0} \frac{dc_0^*}{dW_0} + \beta (W_0 - c_0) \frac{d\alpha_0^*}{dW_0} \underbrace{\mathbb{E}[u'(c_1) r^e]}_{= 0 \text{ by FOC for } \alpha_0} + \beta \mathbb{E}[u'(c_1) r^w]. \end{aligned}$$

Both the first and second terms vanish by the respective first-order conditions, giving the **envelope condition**:

$$V'(W_0) = \beta \mathbb{E}[u'(c_1) r^w] = u'(c_0),$$

where the last equality again uses the FOC for  $c_0$ . The intuition is transparent:  $V'(W_0)$  is the marginal value of a dollar of wealth, while  $u'(c_0)$  is the marginal utility of consuming it today. Their equality is simply the optimality condition — if consuming were worth more than investing, the agent would have consumed more. The middle expression  $\beta \mathbb{E}[u'(c_1) r^w]$  confirms this: the value of investing a dollar is the expected discounted marginal utility next period, scaled by the portfolio return, which at the optimum equals the cost of forgoing a dollar of current consumption.

The derivation assumes  $V$  is differentiable in  $W_0$ , which is not automatic: the value function is defined as a maximum and need not be smooth. Benveniste and Scheinkman (1979) show that differentiability holds when the feasible set is “regular” near the optimum (essentially, when the constraint correspondence has an interior). Milgrom and Segal (2002) extend the result to much more general settings, requiring only that the objective is absolutely continuous in the state parameter — a condition satisfied in standard portfolio problems.

Using the envelope condition  $V'(W_0) = u'(c_0)$ , the SDF can also be written as

$$m = \beta \frac{u'(c_1)}{u'(c_0)} = \beta \frac{u'(W_1)}{V'(W_0)}.$$

The first form is the familiar consumption-based SDF: assets that pay off when  $u'(c_1)$  is high — that is, when future consumption is low and marginal utility is high — command a high price today. The second form expresses the same object using the value function:  $V'(W_0)$  is the shadow price of a dollar of current wealth, and  $u'(W_1)$  is the marginal utility of next-period

wealth. Writing the SDF this way is useful for extending the problem to multiple periods, since  $V'$  plays the role of the “current marginal value” regardless of how many periods lie ahead.

**Example 2** (CRRA Utility). Let  $u(c) = \frac{c^{1-\eta}}{1-\eta}$  so that  $u'(c) = c^{-\eta}$ . The Euler equation for the risk-free asset is

$$\mathbb{E} \left( \beta \left( \frac{c_1}{c_0} \right)^{-\eta} r^f \right) = 1.$$

CRRA utility has no wealth effects on the savings rate: if you double wealth, you optimally double both consumption and investment, keeping the fraction saved constant. This suggests guessing that  $c_0 = kW_0$  for some constant  $k \in (0, 1)$ , i.e., the agent consumes a fixed fraction of wealth. Then  $c_1 = r^w(W_0 - c_0) = r^wW_0(1 - k)$ , so

$$\frac{c_1}{c_0} = \frac{1 - k}{k} r^w.$$

Substituting into the Euler equation:

$$\beta \left( \frac{1 - k}{k} \right)^{-\eta} r^f \mathbb{E}((r^w)^{-\eta}) = 1.$$

Isolating the  $k$ -dependent term and taking both sides to the power  $1/\eta$ :

$$\frac{1 - k}{k} = (\beta r^f \mathbb{E}((r^w)^{-\eta}))^{1/\eta}.$$

Letting  $\phi = (\beta r^f \mathbb{E}((r^w)^{-\eta}))^{1/\eta}$  and solving  $\frac{1-k}{k} = \phi$  gives

$$k = \frac{1}{1 + \phi} = \frac{1}{1 + (\beta r^f \mathbb{E}((r^w)^{-\eta}))^{1/\eta}} < 1.$$

The value function inherits the power form:

$$V(W_0) = \frac{c_0^{1-\eta}}{1-\eta} + \beta \mathbb{E} \left( \frac{c_1^{1-\eta}}{1-\eta} \right) = a \frac{W_0^{1-\eta}}{1-\eta},$$

where

$$a = k^{1-\eta} + \beta(1 - k)^{1-\eta} E((r^w)^{1-\eta}).$$

## References

Benveniste, Lawrence M., and Jose A. Scheinkman. 1979. "On the Differentiability of the Value Function in Dynamic Models of Economics." *Econometrica* 47 (3): 727–32.

Milgrom, Paul, and Ilya Segal. 2002. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica* 70 (2): 583–601.