Consumption and Portfolio Choice

Lorenzo Naranjo

March 2026

Introduction

The stochastic discount factor notebook established that any valid pricing rule is represented by a random variable m satisfying $\mathrm{E}(m\,\mathbf{R}) = \boldsymbol{\iota}$, where $\mathbf{R}$ is the vector of gross returns and $\boldsymbol{\iota}$ is a vector of ones. That derivation was purely mathematical — it required only the law of one price and no-arbitrage, with no reference to preferences or consumption. In Utility Theory Under Uncertainty we studied how agents rank random payoffs through expected utility and characterized their attitudes toward risk through the curvature of u. This notebook puts those preferences to work: we derive the SDF from the optimality conditions of a utility-maximizing investor and show, for the first time, that it equals a ratio of marginal utilities of consumption.

We proceed in two steps, each a dynamic programming problem. In both, the agent's situation is fully summarized by a single **state variable** — initial wealth W_0 — and the **value function** $V(W_0)$ records the maximum expected utility attainable from that state.

We begin with the simplest case — portfolio choice with no current consumption — where the portfolio $\boldsymbol{\alpha}$ is the only control and the Bellman equation collapses to $V(W_0) = \max_{\boldsymbol{\alpha}} \mathrm{E}[u(R^w W_0)]$. The first-order condition immediately expresses the SDF in terms of marginal utility of end-of-period wealth. We then add a second control — current consumption c_0 — obtaining a two-date problem in which the agent trades off consuming today against investing for tomorrow. The **envelope theorem** is the key tool for this second step: differentiating V with respect to W_0 at the optimum eliminates terms involving how the controls respond to the state, yielding $V'(W_0) = u'(c_0)$. This links the marginal value of wealth to the marginal utility of consumption and gives the SDF its second economic interpretation as a ratio of marginal utilities across dates. The intertemporal portfolio choice notebook extends this argument to a general multi-period setting.

Portfolio Choice with No Current Consumption

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

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

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

where $\mathbf{R}^e = \mathbf{R} - R^f\boldsymbol{\iota}$ 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:

$$V(W_0) = \max_{\boldsymbol{\alpha}} \ \mathrm{E}\left(u(R^w W_0)\right).$$

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

$$\mathrm{E}\left(u'(W)\,\mathbf{R}^e\right) = \mathbf{0}.$$

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

$$\mathrm{E}\left(u'(W)\,\mathbf{R}\right) = \mathrm{E}\left(u'(W)\right) R^f\boldsymbol{\iota},$$

which can be rewritten as

$$\mathrm{E}\left(\frac{u'(W)}{R^f\,\mathrm{E}(u'(W))}\,\mathbf{R}\right) = \boldsymbol{\iota}.$$

The first-order condition gives the SDF its first economic interpretation: the pricing kernel is proportional to the investor's marginal utility of end-of-period wealth. Specifically,

$$m = \frac{u'(W)}{R^f\ \mathrm{E}\left(u'(W)\right)}$$

satisfies $\mathsf{E}(m\,\mathbf{R}) = \boldsymbol{\iota}$ directly by the FOC, and $\mathsf{E}(m\,R^f) = 1$ since R^f is deterministic, confirming it is a valid SDF. States where wealth is low — and therefore marginal utility $u'(W)$ is high — receive a large discount factor, making assets that pay off in those states more valuable.

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

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

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

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

CARA Utility

Optimal Portfolio

We now specialize the previous analysis to CARA utility combined with normally distributed returns, a classical setting that yields a closed-form optimal portfolio and an explicit SDF. Let $u(c) = -e^{-\gamma c}$ with $\gamma > 0$, so that $u'(W) = \gamma e^{-\gamma W}$ and the SDF from above is

$$
m = \frac{e^{-\gamma W}}{R^f\ \mathsf{E}(e^{-\gamma W})}. \tag{1}
$$

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 $W_0 = y^d + \mathbf{y}'\boldsymbol{\iota}$ 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

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

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

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

The agent invests a dollar amount proportional to $\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}^e$ in each risky asset. Strikingly, $\mathbf{y}$ does not depend on W_0 — CARA preferences exhibit no wealth effect on risky asset demand. Since $y^d = W_0 - \mathbf{y}'\boldsymbol{\iota}$, any additional wealth flows entirely into the risk-free asset: the risky dollar positions are fixed, and the portfolio weight

$$= \frac{\mathbf{y}}{W_0} = \frac{1}{\gamma W_0}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}^e$$

shrinks as wealth grows. Expressed in weights, this takes the standard mean-variance form with relative risk aversion γW_0, reflecting the fact that CARA has constant absolute risk aversion γ but relative risk aversion that increases linearly with wealth.

This demand function — linear in expected excess returns and inversely proportional to risk aversion — is the building block of demand-based asset pricing, where equilibrium prices are determined by aggregating investor demands across heterogeneous investors (Koijen and Yogo 2019). Moreover, note that $\boldsymbol{\Sigma} = \mathrm{Cov}(\mathbf{R}, \mathbf{R}')$, so that

$$\boldsymbol{\Sigma}\mathbf{y} = \mathrm{Cov}(\mathbf{R}, \mathbf{R}'\mathbf{y}) = \mathrm{Cov}(\mathbf{R}, \mathbf{y}'\mathbf{R}) = \mathrm{Cov}(\mathbf{R}, W),$$

where the last equality uses $W = y^d R^f + \mathbf{y}'\mathbf{R}$ and the fact that $y^d R^f$ is non-random. The excess return vector then satisfies

$$\boldsymbol{\mu}^e = \gamma\boldsymbol{\Sigma}\mathbf{y} = \gamma\,\mathrm{Cov}(\mathbf{R}, W). \tag{2}$$

Thus the risk premium on each asset is proportional to its covariance with end-of-period wealth W.

Stochastic Discount Factor

This result is consistent with the SDF formula derived above. Equation (1) shows that m is a decreasing exponential in wealth; it is high precisely in the states where wealth is low and marginal utility is high. To verify it delivers the same risk premia, recall from the stochastic discount factor notebook that

$$\mathsf{E}(R^i) - R^f = -R^f \,\mathsf{Cov}(m, R^i).$$

Since $W = y^d R^f + \mathbf{y}'\mathbf{R}$ is a linear function of the normal vector $\mathbf{R}$, W and R^i are jointly normal, so Stein's lemma gives $\mathsf{Cov}(e^{-\gamma W}, R^i) = -\gamma\ \mathsf{E}(e^{-\gamma W})\ \mathsf{Cov}(W, R^i)$. Substituting:

$$-R^f\,\mathsf{Cov}(m, R^i) = -R^f \frac{-\gamma\ \mathsf{E}(e^{-\gamma W})\ \mathsf{Cov}(W, R^i)}{R^f\ \mathsf{E}(e^{-\gamma W})} = \gamma\ \mathsf{Cov}(W, R^i).$$

Therefore,

$$\mathsf{E}(R^i) - R^f = \gamma\ \mathsf{Cov}(W, R^i),$$

exactly as derived from the FOC.

CAPM Derivation

We can use the result above to derive the CAPM. In a previous notebook, we derived the CAPM using the beta-pricing relationship together with the assumption that investors have mean-variance preferences. CARA utility with normally distributed returns implies exactly that portfolio-choice problem.

Assume that investors $h = 1, \ldots, H$ all have CARA utility with risk aversion γ_h, face the same normal return distribution, and can borrow or lend at the same risk-free rate. Investor h then demands

$$\mathbf{y}_h = \frac{1}{\gamma_h}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}^e.$$

Aggregating across investors gives

$$\mathbf{y} = \sum_{h=1}^{H} \mathbf{y}_h = \left(\sum_{h=1}^{H} \frac{1}{\gamma_h} \right) \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^e = \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^e,$$

where $\frac{1}{\gamma} = \sum_{h=1}^{H} \frac{1}{\gamma_h}$ and $\mathbf{y}$ is the vector of aggregate risky holdings. If the risk-free asset is in zero net supply, aggregate market wealth is simply

$$W = \mathbf{y}'\mathbf{R}.$$

Under this equilibrium assumption, aggregate wealth is just the payoff on the market portfolio. Let

$$R^M = \frac{W}{\mathbf{y}'\boldsymbol{\iota}}$$

denote the return on the market portfolio. Applying (2) to the aggregate economy therefore gives

$$\mathrm{E}(R^i) - R^f = \gamma \,\mathrm{Cov}(R^i, W) = \gamma\, \mathbf{y}'\boldsymbol{\iota}\; \mathrm{Cov}(R^i, R^M).$$

Setting $i = M$ gives

$$\mathrm{E}(R^M) - R^f = \gamma\, \mathbf{y}'\boldsymbol{\iota}\; \mathrm{V}(R^M),$$

so

$$\gamma\, \mathbf{y}'\boldsymbol{\iota} = \frac{\mathrm{E}(R^M) - R^f}{\mathrm{V}(R^M)}.$$

Substituting back,

$$\mathrm{E}(R^i) - R^f = \frac{\mathrm{Cov}(R^i, R^M)}{\mathrm{V}(R^M)} \left(\mathrm{E}(R^M) - R^f \right) = \beta_{iM} \left(\mathrm{E}(R^M) - R^f \right),$$

where $\beta_{iM} = \mathrm{Cov}(R^i, R^M) / \mathrm{V}(R^M)$. Thus market clearing turns the CARA covariance-pricing relation into the CAPM.

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

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

subject to

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

First-Order Conditions

The FOC with respect to c_0 and $\boldsymbol{\alpha}_0$ are, respectively,

$$u'(c_0) + \beta \; \mathsf{E}\left(u'(c_1) \cdot (-R^w)\right) = 0,$$

$$\beta \; \mathsf{E}\left(u'(c_1)\,(\mathbf{R} - R^f \boldsymbol{\iota})\right) = \mathbf{0}.$$

Combining the two conditions we can rearrange to obtain

$$\mathsf{E}\left(\beta \frac{u'(c_1)}{u'(c_0)}\,\mathbf{R}\right) = \boldsymbol{\iota},$$

which shows that

$$m = \beta \frac{u'(c_1)}{u'(c_0)} \in \mathcal{M} \tag{3}$$

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 $\boldsymbol{\alpha}_0^*$ 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

$$V(W_0) = u\left(c_0^*(W_0)\right) + \beta\ \mathsf{E}\left(u\left((\boldsymbol{\alpha}_0^*(W_0)'\mathbf{R}^e + R^f)(W_0 - c_0^*(W_0))\right)\right).$$

Differentiating with respect to W_0 and using the chain rule:

$$\begin{aligned} V'(W_0) &= u'(c_0)\frac{dc_0^*}{dW_0} + \beta\ \mathsf{E}\left[u'(c_1)\left(R^w\left(1 - \frac{dc_0^*}{dW_0}\right) + (W_0 - c_0)\frac{d\boldsymbol{\alpha}_0^{*\prime}}{dW_0}\mathbf{R}^e\right)\right] \\ &= \underbrace{(u'(c_0) - \beta\ \mathsf{E}[u'(c_1)\, R^w])}_{=\,0 \text{ by FOC for } c_0}\frac{dc_0^*}{dW_0} + \beta\,(W_0 - c_0)\frac{d\boldsymbol{\alpha}_0^{*\prime}}{dW_0}\underbrace{\mathsf{E}[u'(c_1)\,\mathbf{R}^e]}_{=\,\mathbf{0} \text{ by FOC for } \boldsymbol{\alpha}_0} + \beta\ \mathsf{E}[u'(c_1)\, R^w]. \end{aligned}$$

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

$$V'(W_0) = \beta\ \mathsf{E}[u'(c_1)\, R^w] = u'(c_0), \tag{4}$$

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 $\beta\ \mathsf{E}[u'(c_1)\, R^w]$ 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 in (4), the SDF in (3) 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.

CRRA Utility

We now specialize the two-date problem to CRRA utility, $u(c) = c^{1-\eta}/(1-\eta)$, so that $u'(c) = c^{-\eta}$. Using (3), the SDF becomes

$$m = \beta \left(\frac{c_1}{c_0}\right)^{-\eta},$$

The goal of this section is threefold: derive the optimal consumption c_0 and next-period consumption c_1 in closed form, express the SDF as an explicit function of the portfolio return R^w, and verify that the value function inherits the power form of u.

Optimal Consumption

CRRA utility is homogeneous of degree $1-\eta$ in consumption: if all wealth scales by a factor, optimal consumption and savings scale by the same factor, leaving the savings rate unchanged. This homogeneity implies that c_0 is a fixed fraction of wealth. We guess $c_0 = kW_0$ for some constant $k \in (0,1)$ and verify that k is indeed constant. Then $c_1 = R^w(W_0 - c_0) = R^w W_0(1 - k)$, so

$$\frac{c_1}{c_0} = \frac{1-k}{k} R^w. \tag{5}$$

To determine k, we combine the two first-order conditions from above. The FOC for c_0 gives $u'(c_0) = \beta \ \mathsf{E}[u'(c_1)\, R^w]$. Expanding $R^w = \boldsymbol{\alpha}_0' \mathbf{R}^e + R^f$ and applying the FOC for $\boldsymbol{\alpha}_0$ ($\mathsf{E}[u'(c_1)\mathbf{R}^e] = \mathbf{0}$) collapses this to the Euler equation for the risk-free return:

$$u'(c_0) = \beta\, R^f \ \mathsf{E}\left[u'(c_1)\right].$$

Substituting $u'(c) = c^{-\eta}$ and the guessed policies:

$$k^{-\eta} = \beta\, R^f \ (1-k)^{-\eta} \ \mathsf{E}\left[(R^w)^{-\eta}\right].$$

The factor $W_0^{-\eta}$ cancels from both sides, confirming that k is indeed a constant. Isolating the k-dependent term and taking both sides to the power $1/\eta$:

$$\frac{1-k}{k} = \left(\beta\, R^f \ \mathsf{E}\left((R^w)^{-\eta}\right)\right)^{1/\eta}.$$

Letting $\phi = \left(\beta\, R^f \ \mathsf{E}((R^w)^{-\eta})\right)^{1/\eta}$ and solving $\frac{1-k}{k} = \phi$ gives

$$k = \frac{1}{1+\phi} = \frac{1}{1 + \left(\beta\, R^f \ \mathsf{E}\left((R^w)^{-\eta}\right)\right)^{1/\eta}} < 1.$$

Stochastic Discount Factor

Equation (5) implies that the SDF can be written as

$$m = \beta \left(\frac{1-k}{k}\right)^{-\eta} (R^w)^{-\eta},$$

a decreasing function of the portfolio return alone — consumption has been substituted out entirely. This is a consequence of the CRRA structure: because the savings rate $1-k$ is constant, consumption growth is proportional to the portfolio return, so the two representations are equivalent. This foreshadows Epstein-Zin preferences (Epstein and Zin 1989), in which the coefficient of relative risk aversion γ and the elasticity of intertemporal substitution ψ are decoupled. The Epstein-Zin SDF takes the general form

$$m = \beta^{\theta} \left(\frac{c_1}{c_0}\right)^{-\theta/\psi} (R^w)^{\theta-1}, \quad \theta = \frac{1-\gamma}{1-1/\psi},$$

which involves c_1/c_0 and R^w separately when $\psi \neq 1/\gamma$. Under CRRA, $\gamma = \eta$ and $\psi = 1/\eta$, giving $\theta = 1$ and collapsing the Epstein-Zin SDF to $m = \beta(c_1/c_0)^{-\eta}$ — exactly the expression above. The proportionality $c_1/c_0 \propto R^w$ then further reduces it to a pure function of the portfolio return, as seen here.

Value Function

The value function inherits the same power form as u. Substituting $c_0 = kW_0$ and $c_1 = R^w(1-k)W_0$:

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

where

$$a = k^{1-\eta} + \beta(1-k)^{1-\eta}\,\mathsf{E}\left((R^w)^{1-\eta}\right).$$

Differentiating gives $V'(W_0) = a\,W_0^{-\eta}$. The envelope theorem requires $V'(W_0) = u'(c_0) = (kW_0)^{-\eta}$, so $a = k^{-\eta}$ at the optimum — a consistency check that the power-form guess is correct.

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.

Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (4): 937–69.

Koijen, Ralph S. J., and Motohiro Yogo. 2019. “A Demand System Approach to Asset Pricing.” *Journal of Political Economy* 127 (4): 1475–515.

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