

# The Stochastic Discount Factor

Lorenzo Naranjo

March 2026

This notebook introduces the **stochastic discount factor** (SDF), one of the central objects in asset pricing. The SDF is a random variable  $m$  that prices every traded payoff through the formula  $p(x) = E(mx)$ . Its importance is that it compresses all of asset pricing into a single object: once  $m$  is known, we can value any traded cash flow, recover the risk-free rate, and decompose expected returns into time-value-of-money and risk-premium components. The same idea also creates a bridge between no-arbitrage and equilibrium theory. In no-arbitrage language, the SDF is the linear functional that prices payoffs consistently across assets; in equilibrium language, it is the marginal value of wealth or consumption across states, as in Lucas (1978) and Breeden (1979). This is why the SDF, or **pricing kernel**, is the common language of modern asset pricing; see also Cochrane (2009).

Rather than taking the pricing formula as given, we derive it from basic restrictions on asset prices. First, the **law of one price** implies that prices must be generated by a linear pricing rule on the space of traded payoffs. Then the **principle of no-arbitrage** strengthens that result by implying the existence of a strictly positive SDF. In complete markets, this strictly positive SDF is unique. The representation used here follows the SDF approach developed in Hansen and Richard (1987), Hansen and Jagannathan (1991), and Hansen and Jagannathan (1997), with the no-arbitrage foundations rooted in work such as Kreps (1981).

## The Set of Traded Payoffs

Let  $L = \mathbb{R}^S$  denote the full space of state-contingent payoffs: each element of  $L$  specifies one payoff in each of the  $S$  states of the world. In this economy, not every payoff in  $L$  is necessarily traded unless the market is *complete*. In a finite-state economy, we identify each payoff with its vector of state realizations in  $\mathbb{R}^S$ . We use the following bracket convention throughout. Square

brackets display the state-by-state realizations of a random variable or any vector indexed by states; for example, if  $u, v \in \mathbb{R}^S$ , then

$$u = \begin{bmatrix} u(1) \\ \vdots \\ u(S) \end{bmatrix}, \quad \text{and} \quad v = \begin{bmatrix} v(1) \\ \vdots \\ v(S) \end{bmatrix},$$

and the inner product of two vectors in  $\mathbb{R}^S$  is written  $u \cdot v = \sum_{s=1}^S u(s) v(s)$ . A prime ( $'$ ) denotes transposition.

Let  $x_1, x_2, \dots, x_N$  denote linearly independent traded payoffs that form a basis for the marketed payoff space

$$X = \text{span}\{x_1, x_2, \dots, x_N\},$$

where  $N = \dim(X) \leq S$ . Here  $\mathbb{R}^N$  is the coordinate space for portfolios and other stacked traded-asset objects. Vectors in this asset space will typically be written with decorated symbols such as  $\mathbf{x}$  or  $\boldsymbol{\pi}$ ; when the dimension is clear from context, we use the same dot-product notation in either space. Round parentheses stack objects of the same kind into a structural column vector in asset space. Denote by

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{pmatrix} \tag{1}$$

a column vector containing all the basis payoffs. The associated Gram matrix

$$E(\mathbf{x}\mathbf{x}') = \sum_{s=1}^S q(s)\mathbf{x}(s)\mathbf{x}'(s)$$

is invertible because the payoffs  $x_1, \dots, x_N$  are linearly independent.

We can create other payoffs by buying or selling portfolios of these  $N$  traded payoffs. Any  $x \in X$  can be expressed as:

$$x = \sum_{i=1}^N a_i x_i, \tag{2}$$

for  $a_i \in \mathbb{R}$ ,  $1 \leq i \leq N$ . Finally, we denote by  $\pi_i$  the price of traded payoff  $x_i$  for  $1 \leq i \leq N$ ,

and by

$$\boldsymbol{\pi} = \begin{pmatrix} \pi_1 \\ \pi_2 \\ \vdots \\ \pi_N \end{pmatrix}$$

a column vector containing the prices of the  $N$  basis payoffs.

## The Law of One Price

Now, we would like to know if there is a way to create a pricing functional  $p : X \rightarrow \mathbb{R}$  that gives the price of any traded payoff. Clearly, we have that  $p(x_i) = \pi_i$  for  $1 \leq i \leq N$ , and our intuition tells us that the price of any other traded payoff should be pinned down by the prices of the basis assets used to construct it. The next example shows what goes wrong when this consistency fails.

**Example 1.** Suppose that  $p(x) = 1$  and  $p(y) = 2$ . There is also an asset  $z = 3x + 4y$  such that  $p(z) = 12$ . Is there an arbitrage opportunity?

Of course! We could buy 3 units of  $x$  and 4 units of  $y$  and bundle them as  $z$ . The cost of the bundle is \$11, but we can sell it for \$12, generating a riskless profit of \$1 per trade. Since there is no shortage of these securities, we could continue doing it until the prices of  $x$  and  $y$  go up and/or the price of  $z$  goes down.  $\square$

In order to avoid this type of situation, we will assume the following.

**Assumption 1** (The Law of One Price). *Suppose that  $x_i \in X$  and  $a_i \in \mathbb{R}$  for  $i \in \{1, 2, \dots, N\}$ .*

*If  $x = \sum_{i=1}^N a_i x_i \in X$ , then*

$$p(x) = \sum_{i=1}^N a_i p(x_i). \quad (3)$$

In competitive markets, the **law of one price** (LOOP) guarantees that the price of a basket of stocks is equal to the sum of the prices of its constituents. This logic is at the heart of how **Exchange-Traded Funds** (ETF) operate, as the next example shows.

**Example 2.** An ETF is a type of investment fund that is traded on stock exchanges, similar to individual stocks. ETFs hold a diversified portfolio of assets, such as stocks, bonds, or commodities, which provides investors with broad exposure to specific markets or investment strategies.

ETF arbitrage is the mechanism that helps keep the market price of an ETF in line with its Net Asset Value (NAV). Authorized Participants (APs), typically large financial institutions, have the ability to create or redeem ETF shares in large blocks called creation units.

When the ETF market price is higher than the NAV, APs can buy the underlying securities of the ETF in the open market and then deliver them to the ETF issuer in exchange for new ETF shares. The AP can then sell these ETF shares at the higher market price, making a profit. This buying of underlying securities pushes their prices up, while the selling of new ETF shares pushes the ETF price down, bringing the two prices closer together.

When the ETF market price is lower than the NAV, APs can buy ETF shares in the open market and deliver them to the ETF issuer in exchange for the underlying securities. The AP can then sell these underlying securities at the higher NAV price, making a profit. This buying of ETF shares pushes their price up, while the selling of the underlying securities pushes their prices down, again bringing the two prices closer together.

This creation and redemption process happens continuously and helps to keep the ETF price in line with the NAV. The arbitrage opportunities are typically small but sufficient for APs to engage in the process for profit, ensuring that the ETF price does not deviate significantly from its NAV.

You can find more information [here](#). □

The fact that the market for ETFs is so liquid and works flawlessly reassures us that LOOP is a reasonable axiom to start working from. The law of one price implies that the price functional defined in (3) is linear on the finite-dimensional payoff space  $X$ , hence continuous. By the [Riesz Representation Theorem](#), there exists a unique  $x^* \in X$  such that  $p(x) = E(x^*x)$  for all  $x \in X$ .

Since  $p(x_i) = \pi_i$  for each basis asset, the representing payoff  $x^*$  itself satisfies  $E(x^*\mathbf{x}) = \boldsymbol{\pi}$  and is therefore a valid SDF. More generally, if  $m$  is any valid SDF, then  $E(m\mathbf{x}) = \boldsymbol{\pi}$ . Substituting

into the [Projection onto a Subspace](#) formula (with  $y = m$ ,  $M = X$ , and  $G = E(\mathbf{x}\mathbf{x}')$ ) identifies  $x^*$  as the projection of any valid SDF onto  $X$ .

**Property 1.** Under LOOP, the unique SDF in  $X$  is the projection of any valid SDF  $m$  onto  $X$ , given by

$$x^* = \boldsymbol{\pi}' E(\mathbf{x}\mathbf{x}')^{-1} \mathbf{x} \in X. \quad (4)$$

To verify directly, take any  $x = \mathbf{x}' \mathbf{a} \in X$ , where  $\mathbf{a} = (a_1, \dots, a_N)'$ , and compute

$$\begin{aligned} E(x^* x) &= E(\boldsymbol{\pi}' E(\mathbf{x}\mathbf{x}')^{-1} \mathbf{x} \mathbf{x}' \mathbf{a}) \\ &= \boldsymbol{\pi}' E(\mathbf{x}\mathbf{x}')^{-1} E(\mathbf{x}\mathbf{x}') \mathbf{a} \\ &= \boldsymbol{\pi}' \mathbf{a} = p(x), \end{aligned}$$

confirming that  $x^*$  prices all assets correctly.

**Example 3.** As a concrete example, suppose there are four states with equal probabilities  $P(\omega) = 0.25$  and two traded payoffs

$$x_1 = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \quad x_2 = \begin{bmatrix} 0 \\ 1 \\ 1 \\ 2 \end{bmatrix}.$$

Collect them in the column vector

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

So  $X = \text{span}\{x_1, x_2\}$ . Take a valid SDF

$$m = \begin{bmatrix} 1.2 \\ 0.8 \\ 1.0 \\ 0.6 \end{bmatrix},$$

which need not belong to  $X$ . The corresponding asset prices are

$$\boldsymbol{\pi} = E(m\mathbf{x}) = \begin{pmatrix} E(mx_1) \\ E(mx_2) \end{pmatrix} = \begin{pmatrix} 0.5 \\ 0.75 \end{pmatrix}.$$

Also,

$$E(\mathbf{x}\mathbf{x}') = \begin{pmatrix} E(x_1^2) & E(x_1x_2) \\ E(x_1x_2) & E(x_2^2) \end{pmatrix} = \begin{pmatrix} 0.5 & 0.25 \\ 0.25 & 1.5 \end{pmatrix}.$$

Applying equation (4),

$$x^* = \boldsymbol{\pi}' E(\mathbf{x}\mathbf{x}')^{-1} \mathbf{x} = \frac{9}{11}x_1 + \frac{4}{11}x_2.$$

This  $x^*$  lies in the payoff space  $X$  and satisfies

$$E(x^*x_1) = 0.5, \quad E(x^*x_2) = 0.75,$$

so it prices the traded assets exactly even though the original SDF  $m$  is not itself in  $X$ . □

Any other SDF  $m$  will price the assets correctly, so that

$$E((m - x^*)x) = E(mx) - E(x^*x) = p(x) - p(x) = 0.$$

This shows that we can create new SDFs by combining  $x^*$  with any random variable  $\varepsilon$  orthogonal to  $X$ . In other words, all the SDFs that price assets correctly in  $X$  can be written as:

$$m = x^* + \varepsilon,$$

where  $\varepsilon \perp x$  for all  $x \in X$ .

**Example 4.** Continuing Example 3, the orthogonal component is

$$\varepsilon = m - x^* = \begin{bmatrix} \frac{21}{55} \\ \frac{21}{55} \\ -\frac{7}{55} \\ \frac{11}{7} \\ -\frac{11}{55} \end{bmatrix}.$$

Using the equal state probabilities, we obtain

$$E(\varepsilon x_1) = 0.25 \left( \frac{21}{55} - \frac{21}{55} \right) = 0$$

and

$$E(\varepsilon x_2) = 0.25 \left( -\frac{21}{55} + \frac{7}{11} - \frac{14}{55} \right) = 0.$$

Therefore, for any traded payoff  $x = a_1 x_1 + a_2 x_2 \in X$ ,

$$E(x\varepsilon) = a_1 E(x_1\varepsilon) + a_2 E(x_2\varepsilon) = 0,$$

which verifies in this example that  $\varepsilon \in X^\perp$ . □

## The Principle of No-Arbitrage

The law of one price disciplines prices across portfolios: equal payoffs must have equal prices. But LOOP says nothing about the relationship between prices and the *sign* of payoffs. A pricing functional can be perfectly linear — satisfying LOOP — and yet assign a non-positive price to a payoff that is non-negative in every state. An investor facing such an opportunity would want to hold an unbounded position: buying costs nothing (or generates an immediate cash inflow) while the payoff can only be zero or positive. In competitive markets, unbounded demand is inconsistent with equilibrium clearing, so such opportunities cannot persist.

**Assumption 2** (Principle of No-Arbitrage). *The price of a payoff that is non-negative in all states and strictly positive in at least one state of the world must be positive.*

To see that PNA is strictly stronger than LOOP, consider a market with a single risky asset  $x_1$  that pays \$1 in state 1 and \$0 in state 2. Any pricing functional of the form  $p(ax_1) = a\pi_1$  is linear and therefore satisfies LOOP regardless of the sign of  $\pi_1$ . But if  $\pi_1 \leq 0$ , buying  $x_1$  costs nothing (or generates income) while delivering a non-negative payoff. LOOP is silent on this because there is no inconsistency across portfolios; the problem is that the price assigns the wrong sign to a desirable payoff.

More generally, PNA adds the requirement that prices respect *dominance*: a payoff that weakly improves on zero in every state, and strictly improves in at least one, must carry a strictly positive price. This is the minimal condition consistent with the behavior of an investor who strictly prefers more consumption to less in every state of the world.

The principle of no-arbitrage (PNA) is a strictly stronger assumption than LOOP, as the next property shows.

**Property 2.**

$$PNA \Rightarrow LOOP.$$

Proof

We will prove this claim by contradiction. Assume that PNA holds but not LOOP. A violation of LOOP implies that the price of a zero payoff is not zero: if  $p(x) \neq \sum_i a_i \pi_i$  for some  $x = \sum_i a_i x_i \in X$ , then the zero payoff  $x - \sum_i a_i x_i = 0$  is priced at  $p(x) - \sum_i a_i \pi_i \neq 0$ . Without loss of generality, assume  $p_0 = p(x) - \sum_i a_i \pi_i > 0$ . Take any payoff  $x^+ \in X$  that is non-negative and non-zero, with price  $p > 0$ . Form a portfolio that buys one unit of  $x^+$  and sells  $n > \frac{p}{p_0}$  units of this zero-payoff portfolio. The cost of that portfolio is

$$\pi = p - np_0 < p - \frac{p}{p_0} p_0 = 0,$$

but its payoff is identical to  $x^+$ , which is non-negative and non-zero, violating PNA.

This contradiction establishes the result. □

The reverse implication does not hold: a market can satisfy LOOP while still admitting free lunches, as the single-asset example above illustrates. PNA therefore imposes a genuinely additional restriction on the pricing functional beyond linearity. In particular, it rules out not only mispricing across equivalent portfolios, but also prices that attach non-positive value to payoffs that are everywhere non-negative. This distinction is important because return-based formulas implicitly assume that the asset price is strictly positive. If a non-negative payoff had zero or negative price, its gross return would be undefined or economically ill-behaved. That is why arbitrage restrictions are stated first in payoff-price language, and only afterwards translated into statements about returns.

A more important consequence of PNA is that it is equivalent to the existence of a strictly positive SDF. Under LOOP, an SDF  $m$  prices every traded payoff via  $p(x) = E(mx)$ , but a generic SDF need not be positive in every state. If some component satisfies  $m(s) \leq 0$ , then the state- $s$  payoff receives non-positive marginal value. In complete markets, the Arrow-Debreu security  $e_s$  would then have price  $p(e_s) = E(me_s) = q(s)m(s) \leq 0$ , immediately violating PNA. More generally, the positive-SDF theorem below shows that ruling out such statewise negative valuations on the traded payoff space is exactly equivalent to PNA.

**Property 3** (No-Arbitrage and Strictly Positive SDFs).

$$PNA \Leftrightarrow \exists \text{ a valid SDF } m > 0.$$

A proof is provided in [the appendix](#). Even in a finite state space, this is not an elementary result: the argument already requires a geometric reformulation of the pricing problem and an application of the Separating Hyperplane Theorem. That depth is part of the point. The equivalence says that no-arbitrage is exactly the condition that the market assigns a positive shadow value to wealth in every state — no state is “throwaway.” In complete markets ( $N = S$ ), PNA pins down a *unique* strictly positive SDF, since  $X = L$  leaves no room for the orthogonal residual  $\varepsilon$  in the decomposition  $m = x^* + \varepsilon$ .

## Arrow-Debreu Securities

Arrow-Debreu securities are the most elementary financial contracts imaginable: security  $e_s$  pays \$1 if state  $s$  occurs and \$0 in every other state. They are the atomic building blocks of the full payoff space  $L$ : every contingent claim can be decomposed state by state using these elementary random variables, whether or not the corresponding securities are actually traded.

Formally, define the random variable  $e_s$  by

$$e_s(j) = \begin{cases} 1 & \text{if } j = s, \\ 0 & \text{otherwise,} \end{cases}$$

for each  $s \in \{1, 2, \dots, S\}$ . Since the payoff space  $L$  has dimension  $S$ , the collection  $\{e_1, e_2, \dots, e_S\}$  forms a basis for  $L$ . Any payoff  $x \in L$  therefore admits the unique decomposition

$$x = x(1) e_1 + x(2) e_2 + \dots + x(S) e_S, \quad (5)$$

since in state  $s$  the right-hand side evaluates to  $x(s) \mathbf{1} = x(s)$ .

This decomposition is purely geometric. It does not mean that every Arrow-Debreu security is traded. When  $N < S$ , the marketed payoff space  $X$  is only a subspace of  $L$ , so many contingent claims in  $L$  need not be attainable by portfolios of traded assets.

If an Arrow-Debreu security  $e_s$  is traded, its price is called the **state price** for state  $s$  and is denoted  $\psi(s)$ . In a complete market, where all  $S$  Arrow-Debreu securities are traded, the price of any payoff  $x \in L$  follows immediately from the spanning representation:

$$p(x) = \sum_{s=1}^S \psi(s) x(s) = \psi \cdot x,$$

where

$$\psi = \begin{bmatrix} \psi(1) \\ \vdots \\ \psi(S) \end{bmatrix}$$

collects all state prices. In that case prices are inner products with the state-price vector.

More generally, even when Arrow-Debreu securities are not traded individually, any SDF  $m$  induces a *state-price vector*

$$\psi(s) = q(s) m(s),$$

and the pricing equation becomes

$$p(x) = E(mx) = \sum_{s=1}^S q(s) m(s) x(s) = \sum_{s=1}^S \psi(s) x(s) = \psi \cdot x.$$

Thus an SDF and a state-price vector are just two representations of the same linear pricing rule: the SDF  $m$  weights payoffs by their probability-adjusted marginal value, while the state-price vector  $\psi$  absorbs the probabilities directly, so that  $p(x) = E(mx) = \psi \cdot x$  holds in either form.

The two are linked by the change of variables  $\psi(s) = q(s) m(s)$ , which is invertible whenever  $q(s) > 0$ . The positivity condition  $m > 0$  is therefore equivalent to  $\psi(s) > 0$  for all  $s$ .

**Example 5.** Continuing Example 3, the original SDF

$$m = \begin{bmatrix} 1.2 \\ 0.8 \\ 1.0 \\ 0.6 \end{bmatrix}$$

induces the state-price vector

$$\psi^m = \frac{1}{4}m = \begin{bmatrix} 0.3 \\ 0.2 \\ 0.25 \\ 0.15 \end{bmatrix}.$$

The projected SDF

$$x^* = \frac{9}{11}x_1 + \frac{4}{11}x_2$$

induces instead

$$\psi^* = \frac{1}{4}x^* = \begin{bmatrix} \frac{9}{44} \\ \frac{13}{44} \\ \frac{1}{11} \\ \frac{2}{11} \end{bmatrix}.$$

These two state-price vectors are different because the market is incomplete, but they agree on traded assets:

$$p(x_1) = \psi^m \cdot x_1 = 0.3 + 0.2 = 0.5,$$

$$p(x_2) = \psi^m \cdot x_2 = 0.2 + 0.25 + 2(0.15) = 0.75,$$

and

$$p(x_1) = \psi^* \cdot x_1 = \frac{9}{44} + \frac{13}{44} = 0.5,$$

$$p(x_2) = \psi^* \cdot x_2 = \frac{13}{44} + \frac{1}{11} + 2\left(\frac{2}{11}\right) = 0.75.$$

So both  $\psi^m$  and  $\psi^*$  represent valid pricing rules for the traded payoff space  $X$ . □

If all  $S$  Arrow-Debreu securities are traded — that is,  $N = S$  and  $X = L$  — the market is

**complete.** In this case the state-price vector  $\psi$  is uniquely determined by the  $S$  asset prices, and so is the SDF:  $m(s) = \psi(s)/q(s)$  for each  $s$ . Incomplete markets ( $N < S$ ) admit multiple valid state-price vectors and hence multiple SDFs, all of which price the  $N$  traded assets correctly but assign different values to non-traded payoffs.

## Returns and the Risk-Free Rate

We now translate the pricing relation into a restriction on expected returns. For any traded payoff  $x$  with non-zero price  $p(x) \neq 0$ , define its gross return by  $R = x/p(x)$ . We work with gross returns throughout this section.

Since  $p(x) = E(mx)$ , dividing both sides by  $p(x)$  gives

$$1 = E(mR).$$

Thus, every traded return has discounted expectation equal to one under any valid SDF.

Assume now there is a traded risk-free asset that pays 1 in every state of the world. Its price satisfies

$$p(1) = E(m),$$

for any valid SDF  $m$ . Therefore  $E(m)$  is the same for all SDFs, because they must all assign the same price to the same traded payoff. PNA implies  $p(1) > 0$ . The gross return of this asset, which is constant in all states of the world, is the risk-free rate

$$R^f = \frac{1}{p(1)} = \frac{1}{E(m)}.$$

Thus, the risk-free rate is the reciprocal of the expected SDF. A higher average discount factor implies a lower risk-free rate.

For any traded asset  $i$ ,

$$1 = E(mR^i) = E(m)E(R^i) + \text{Cov}(m, R^i).$$

Substituting  $E(m) = 1/R^f$  and rearranging:

$$E(R^i) - R^f = -\frac{\text{Cov}(m, R^i)}{E(m)} = -R^f \text{Cov}(m, R^i). \quad (6)$$

The excess return of any asset equals  $-R^f$  times its covariance with the SDF. Assets that tend to pay off when the SDF is high — that is, when the marginal value of wealth is high, typically in bad times — are valuable as insurance and therefore command lower excess returns. Assets that pay off when the SDF is low, in good times, offer little insurance and must compensate investors with higher expected returns.

## The Hansen-Jagannathan Bound

The **Sharpe ratio** of asset  $i$  is its excess expected return per unit of return volatility:

$$SR^i = \frac{E(R^i) - R^f}{\sigma(R^i)},$$

provided  $\sigma(R^i) > 0$ . It is the most widely used measure of risk-adjusted performance. We will show that any valid SDF places a universal upper bound on the Sharpe ratio of every traded asset.

Writing the covariance in equation (6) in terms of the correlation  $\rho_{m,i}$ ,

$$\begin{aligned} E(R^i) - R^f &= -\frac{\text{Cov}(m, R^i)}{E(m)} \\ &= -\frac{\rho_{m,i} \sigma(m) \sigma(R^i)}{E(m)}. \end{aligned}$$

Dividing by  $\sigma(R^i)$  and taking absolute values gives

$$\left| \frac{E(R^i) - R^f}{\sigma(R^i)} \right| = |\rho_{m,i}| \frac{\sigma(m)}{E(m)}.$$

Since the [Cauchy-Schwarz inequality](#) implies  $|\rho_{m,i}| \leq 1$ , we obtain

$$\left| \frac{E(R^i) - R^f}{\sigma(R^i)} \right| \leq \frac{\sigma(m)}{E(m)}. \quad (7)$$

This is the **Hansen-Jagannathan bound** (Hansen and Jagannathan 1991): the absolute Sharpe ratio of any traded asset is bounded above by the coefficient of variation  $\sigma(m)/E(m)$  of the SDF. The bound holds for every valid SDF  $m$  and every traded asset  $i$  with non-zero return volatility. Equality holds when  $|\rho_{m,i}| = 1$ , i.e., when the return is perfectly correlated with the SDF up to sign.

Taking the supremum over all traded returns with non-zero volatility gives

$$\sup_{R \in \mathcal{R}} \left| \frac{E(R) - R^f}{\sigma(R)} \right| \leq \frac{\sigma(m)}{E(m)},$$

where  $\mathcal{R}$  is the set of traded gross returns. Since this inequality holds for every valid SDF, taking the infimum over the set  $\mathcal{M}$  of valid SDFs yields

$$\sup_{R \in \mathcal{R}} \left| \frac{E(R) - R^f}{\sigma(R)} \right| \leq \inf_{m \in \mathcal{M}} \frac{\sigma(m)}{E(m)}.$$

In fact, Cauchy-Schwarz is sharp, so

$$\sup_{R \in \mathcal{R}} \left| \frac{E(R) - R^f}{\sigma(R)} \right| = \inf_{m \in \mathcal{M}} \frac{\sigma(m)}{E(m)}.$$

This equality remains valid in infinite-dimensional Hilbert spaces, although the supremum or infimum need not be attained. It says that the largest feasible Sharpe ratio is exactly the smallest coefficient of variation among all valid SDFs.

The economic content is immediate. Since  $\sigma(m)/E(m) = R^f \sigma(m)$ , a large observed Sharpe ratio forces every valid SDF to be sufficiently volatile relative to its mean. Historically, the U.S. equity market has delivered a Sharpe ratio of roughly 0.5 per year, so any admissible SDF must satisfy  $\sigma(m)/E(m) \geq 0.5$ . Standard consumption-based models with power utility and observed aggregate consumption growth typically generate SDFs that are too smooth to meet this restriction unless risk aversion is implausibly high. This is the **equity premium puzzle** of

Mehra and Prescott (1985), and the Hansen-Jagannathan bound gives its sharpest quantitative expression.

## Appendix

### Proof of Property 3

The proof relies on the following geometric result.

**Lemma 1** (Separating Hyperplane Theorem). *If  $A$  and  $B$  are non-empty disjoint convex subsets of  $\mathbb{R}^n$ , there exists a non-zero vector  $\phi \in \mathbb{R}^n$  and a scalar  $\alpha \in \mathbb{R}$  such that*

$$\phi \cdot a \leq \alpha \leq \phi \cdot b \quad \text{for all } a \in A, b \in B.$$

See, for example, Boyd and Vandenberghe (2004).

*Road map.* The direction ( $\Leftarrow$ ) is direct: if  $m > 0$ , every non-negative non-zero payoff receives a positive price, ruling out arbitrage. For ( $\Rightarrow$ ), we argue by contradiction. We reformulate the problem as finding a strictly positive state-price vector that correctly prices all assets. Assume no such vector exists, so the plane of valid state prices and the positive orthant are disjoint. The Separating Hyperplane Theorem then places a hyperplane between them, producing a separating normal vector. We show this normal is the payoff of some traded portfolio, and that the three properties forced by the separation — non-negative payoff, non-zero payoff, and non-positive cost — make that portfolio an arbitrage, contradicting PNA.

( $\Leftarrow$ ): Suppose  $m > 0$  and let  $x \in X$  satisfy  $x \geq 0$  and  $x \neq 0$ . Since  $x$  is non-negative but not identically zero, there is at least one state  $s^*$  with  $x(s^*) > 0$ . In the sum

$$p(x) = E(mx) = \sum_{s=1}^S q(s) m(s) x(s),$$

every term is non-negative (because  $q(s) > 0$ ,  $m(s) > 0$ , and  $x(s) \geq 0$ ), and the term at  $s = s^*$  is strictly positive. Hence  $p(x) > 0$ , which is exactly PNA.

( $\Rightarrow$ ): Suppose PNA holds. Since PNA  $\Rightarrow$  LOOP (by Property 2), the SDF  $x^*$  exists by Property 1. We must show that some SDF is strictly positive in every state.

*Reformulation in terms of state prices.* Recall that the state-price vector  $\psi \in \mathbb{R}^S$  is given by

$$\psi(s) = q(s) m(s).$$

Since  $q(s) > 0$ , we have  $m(s) > 0$  if and only if  $\psi(s) > 0$ , so proving  $m > 0$  is equivalent to finding  $\psi \gg 0$ . In terms of  $\psi$ , the pricing condition  $\pi_i = E(mx_i)$  becomes a Euclidean dot product  $\pi_i = \psi \cdot x_i$ . Stacking the  $N$  conditions into the  $S \times N$  payoff matrix  $D$  (whose  $i$ -th column is  $x_i$ ),

$$D = \begin{bmatrix} x_1(1) & \cdots & x_N(1) \\ \vdots & \ddots & \vdots \\ x_1(S) & \cdots & x_N(S) \end{bmatrix},$$

the pricing constraints become  $D'\psi = \boldsymbol{\pi}$ . Note that the traded payoff space is  $X = \text{col}(D) = \{D\boldsymbol{\theta} : \boldsymbol{\theta} \in \mathbb{R}^N\}$ , so its orthogonal complement is  $X^\perp = \{e \in \mathbb{R}^S : D'e = 0\}$ . We therefore need to show that  $D'\psi = \boldsymbol{\pi}$  has a solution with  $\psi \gg 0$ . Equivalently, if  $\mathbf{x}$  denotes the column vector of traded payoff random variables introduced earlier, then the  $s$ -th row of  $D$  is  $\mathbf{x}(s)'$ , so  $D$  is the state-by-state realization matrix of  $\mathbf{x}$ .

*Geometric setup.* Define

$$\Psi = \{\psi \in \mathbb{R}^S : D'\psi = \boldsymbol{\pi}\},$$

the flat plane of all state-price vectors consistent with observed prices. It is non-empty because  $\psi^*(s) = q(s) x^*(s) \in \Psi$ . The open positive orthant  $\mathbb{R}_{++}^S$  is the set of strictly positive state prices. Our goal is to show that  $\Psi \cap \mathbb{R}_{++}^S \neq \emptyset$  under PNA.

*Contradiction.* Suppose  $\Psi \cap \mathbb{R}_{++}^S = \emptyset$ . Since both sets are non-empty and convex, the **Separating Hyperplane Theorem** yields a non-zero  $\phi \in \mathbb{R}^S$  and  $\alpha \in \mathbb{R}$  with

$$\phi \cdot \psi \leq \alpha \leq \phi \cdot y \quad \text{for all } \psi \in \Psi, y \in \mathbb{R}_{++}^S.$$

**Step 1:**  $\phi \geq 0$  and  $\alpha \leq 0$ . The positive orthant is a cone: if  $y \in \mathbb{R}_{++}^S$ , then  $ty \in \mathbb{R}_{++}^S$  for every

$t > 0$ . Hence the separation inequality gives

$$\alpha \leq \phi \cdot (ty) = t(\phi \cdot y) \quad \text{for all } t > 0.$$

Letting  $t \rightarrow 0^+$  yields  $\alpha \leq 0$ . Next, suppose for contradiction that  $\phi(s) < 0$  for some state  $s$ . For any  $\varepsilon > 0$  and any  $t > 0$ , the vector

$$y = \varepsilon \mathbf{1} + te_s$$

lies in  $\mathbb{R}_{++}^S$ , so separation implies

$$\alpha \leq \phi \cdot y = \varepsilon \phi \cdot \mathbf{1} + t \phi(s).$$

Since  $\phi(s) < 0$ , the right-hand side tends to  $-\infty$  as  $t \rightarrow \infty$ , contradicting the fixed lower bound  $\alpha$ . Therefore every coordinate satisfies  $\phi(s) \geq 0$ , i.e.,  $\phi \geq 0$ .

**Step 2:  $\phi \in X$ , so  $\phi$  is the payoff of some portfolio.** Fix any  $\psi_0 \in \Psi$ . Since

$$\Psi = \psi_0 + X^\perp,$$

every vector of the form  $\psi_0 + te$  with  $e \in X^\perp$  and  $t \in \mathbb{R}$  also lies in  $\Psi$ . Therefore

$$\phi \cdot (\psi_0 + tv) = \phi \cdot \psi_0 + t(\phi \cdot v) \leq \alpha \quad \text{for all } v \in X^\perp, t \in \mathbb{R}.$$

A linear function of  $t$  that is bounded above for all  $t \in \mathbb{R}$  must have zero slope, so  $\phi \cdot v = 0$  for every  $v \in X^\perp$ . Hence  $\phi \in (X^\perp)^\perp = X$ , and therefore there exists a portfolio  $\theta \in \mathbb{R}^N$  with  $\phi = D\theta$ . For any  $\psi \in \Psi$ ,

$$\phi \cdot \psi = (D\theta) \cdot \psi = \theta \cdot (D'\psi) = \theta \cdot \pi = \pi \cdot \theta,$$

so  $\phi \cdot \psi$  is constant on  $\Psi$ . Since separation gives  $\phi \cdot \psi \leq \alpha$  for every  $\psi \in \Psi$ , it follows that

$$\pi \cdot \theta \leq \alpha.$$

**Step 3: Arbitrage.** The portfolio  $\theta$  satisfies:

- *Non-negative payoff in every state:*  $D\theta = \phi \geq 0$ .
- *Strictly positive payoff in at least one state:*  $\phi \neq 0$  (since  $\phi$  is a non-zero separating vector).
- *Non-positive cost:*  $\pi \cdot \theta \leq \alpha \leq 0$ .

This is precisely an arbitrage, contradicting PNA.

*Conclusion.* The assumption  $\Psi \cap \mathbb{R}_{++}^S = \emptyset$  leads to an arbitrage, so it must be false. Hence there exists  $\psi \gg 0$  with  $D'\psi = \pi$ , and the SDF  $m(s) = \psi(s)/q(s) > 0$  is strictly positive in every state. □

## References

- Boyd, Stephen, and Lieven Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.
- Breeden, Douglas T. 1979. "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities." *Journal of Financial Economics* 7 (3): 265–96.
- Cochrane, John. 2009. *Asset Pricing: Revised Edition*. Princeton university press.
- Hansen, Lars Peter, and Ravi Jagannathan. 1991. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy* 99 (2): 225–62.
- Hansen, Lars Peter, and Ravi Jagannathan. 1997. "Assessing Specification Errors in Stochastic Discount Factor Models." *Journal of Finance* 52 (2): 557–90.
- Hansen, Lars Peter, and Scott F Richard. 1987. "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models." *Econometrica* 55 (3): 587–613.
- Kreps, David M. 1981. "Arbitrage and Equilibrium in Economies with Infinitely Many Commodities." *Journal of Mathematical Economics* 8 (1): 15–35.

Lucas, Robert E. 1978. "Asset Prices in an Exchange Economy." *Econometrica* 46 (6): 1429–45.

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15 (2): 145–61.