The Geometry of the Payoff Space

Lorenzo Naranjo

March 2026

This notebook develops the geometric machinery underlying the theory of the stochastic discount factor. We work throughout in a finite-dimensional probability space, where every random variable is simply a vector of state-contingent payoffs and expectation defines an inner product. This inner-product structure turns the space of payoffs into a Euclidean space, making geometric tools available for economic problems.

Three results organize the notebook. First, we characterize **projections**: the projection of a payoff y onto a subspace M is the element of M closest to y, and its residual is orthogonal to every element of M. This is the geometric content of linear regression. Second, we establish the **Cauchy-Schwarz inequality**, which bounds the correlation between any two payoffs and underlies the Hansen-Jagannathan volatility bound. Third, we prove the **Riesz representation theorem**: every linear pricing functional on the payoff space can be written as an inner product with a unique random variable. That random variable is the stochastic discount factor.

Probability Structure

Uncertainty is represented by a finite set $\mathcal{S} = \{1, \ldots, S\}$ of states, defining a finite probability space $(\mathcal{S}, q)$. The set of all random variables defined in $\mathcal{S}$ is denoted by L and is called the **payoff space**. Thus, for any $x \in L$ we have that

$$x = \begin{bmatrix} x(1) \\ x(2) \\ \vdots \\ x(S) \end{bmatrix} \in \mathbb{R}^S$$

defines all the possible payoffs in each state, and the probability of getting a payoff in a particular state is given by $\mathrm{P}(x = x(s)) = q(s)$ for all $s \in \mathcal{S}$. We assume throughout that $q(s) > 0$ for all $s \in \mathcal{S}$, that is, we will not consider possible outcomes that happen with probability zero.

The payoff space is clearly a linear vector space since for any $x, y \in L$ and $\alpha, \beta \in \mathbb{R}$ we have that $\alpha x + \beta y \in L$. We endow the payoff space with an inner product $\langle \cdot, \cdot \rangle : L \times L \to \mathbb{R}$ defined such that for any $x, y \in L$, we have that

$$\langle x, y \rangle = \mathrm{E}(xy) = \sum_{s=1}^{S} q(s)x(s)y(s).$$

In finite-dimensional spaces, we can use the inner product to define the Euclidean norm $\|\cdot\| : L \to \mathbb{R}^+$ for all $x \in L$ as

$$\|x\| = \sqrt{\langle x, x \rangle}.$$

Clearly, $\|x\| = 0 \Leftrightarrow x = 0$. The second moment of x, defined as $\|x\|^2 = \mathrm{E}(x^2)$, measures the mean squared payoff across states and serves as the natural notion of distance between payoffs: $\|y - z\|$ is small when y and z are close in every state on average.

Note that the inner product can be related to standard statistical moments. For any $x, y \in L$, we have

$$\langle x, y \rangle = \mathrm{E}(xy) = \mathrm{Cov}(x, y) + \mathrm{E}(x)\,\mathrm{E}(y),$$

where $\mathrm{Cov}(x, y) = \mathrm{E}[(x - \mathrm{E}(x))(y - \mathrm{E}(y))]$ denotes the covariance between x and y. In particular, $\|x\|^2 = \mathrm{V}(x) + [\mathrm{E}(x)]^2$, where $\mathrm{V}(x) = \mathrm{Cov}(x, x)$ is the variance of x.

Since L consists of all random variables on $\mathcal{S}$, it is a vector space of dimension S. A natural basis is provided by the *Arrow-Debreu securities* $\{e_1, e_2, \dots, e_S\}$, where

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

for each $s \in \mathcal{S}$. The security e_s pays one unit in state s and zero in every other state. Any payoff $x \in L$ can be expressed in this basis as

$$x = \sum_{s=1}^{S} x(s)\, e_s,$$

since in state i the right-hand side evaluates to $x(i)$. Under the inner product defined above, the Arrow-Debreu securities are mutually orthogonal: for $s \neq t$, the product $e_s e_t$ is zero in every state, so $\langle e_s, e_t \rangle = \mathrm{E}(e_s e_t) = 0$. Moreover, $\langle e_s, e_s \rangle = \mathrm{E}(e_s^2) = q(s)$, so the Arrow-Debreu securities form an *orthogonal*—though not orthonormal—basis for L.

Projections

Given $x, y \in L$, consider the vectors $y_x = \alpha x$ and $z = y - y_x$. We say that y_x is the projection of y on the subspace generated by $\{x\}$ if the norm of z is minimal. To obtain the projection, we need to compute the α that minimizes $\|z\|^2 = \|y - \alpha x\|^2 = \mathrm{E}[(y - \alpha x)^2]$. The first-order condition of this problem is:

$$0 = \mathrm{E}[(y - \alpha x)x] = \langle y - \alpha x, x \rangle = \langle z, x \rangle,$$

which implies that $\alpha = \dfrac{\langle x, y \rangle}{\langle x, x \rangle}$ and $\langle z, y_x \rangle = 0$. Statistically, α is the population regression coefficient of y on x with no intercept: it is the unique scalar that makes the residual $z = y - \alpha x$ uncorrelated with x.

We say that two vectors $x, y \in L$ are orthogonal if their inner product is equal to zero. Thus, we have that $y_x \perp z$, implying that the vector y can be decomposed into two orthogonal components. Indeed, we have that

$$\|y\|^2 = \|z + y_x\|^2 = \|z\|^2 + 2\langle z, y_x \rangle + \|y_x\|^2 = \|z\|^2 + \|y_x\|^2,$$

which is a generalization of the classical Pythagorean theorem.

Property 1 (Orthogonal Decomposition). *Given $x, y \in L$, the projection of y on the subspace generated by $\{x\}$ is given by $y_x = \dfrac{\langle x, y \rangle}{\langle x, x \rangle} x$. The vector $z = y - y_x$ is orthogonal to y_x, implying that*

$$\|y\|^2 = \|z\|^2 + \|y_x\|^2. \tag{1}$$

Equation (1) implies that $\|y\|^2 \geq \|y_x\|^2$, with equality occurring whenever y is proportional to x. Therefore, we have that

$$\|y\|^2 \geq \|y_x\|^2 = \left\| \frac{\langle x, y \rangle}{\langle x, x \rangle} x \right\|^2 = \frac{\langle x, y \rangle^2}{\|x\|^2}.$$

The previous expression is known as the Cauchy-Schwartz inequality and is fundamental in the study of Euclidean vector spaces.

Property 2 (Cauchy-Schwartz Inequality)**.** *Given* $x, y \in L$ *we have that*

$$|\langle x, y\rangle| \leq \|x\|\|y\|. \tag{2}$$

The Cauchy-Schwartz inequality implies that, for any two nonzero $x, y \in L$, the ratio $\dfrac{\langle x, y\rangle}{\|x\|\|y\|}$ lies in $[-1, 1]$ and can be interpreted as the cosine of the angle $\theta \in [0, \pi]$ between x and y. Specifically, $|\langle x, y\rangle| = \|x\|\|y\|$ if and only if y is proportional to x, i.e., the two vectors are collinear. When $\mathsf{E}(x) = \mathsf{E}(y) = 0$, the inner product reduces to $\langle x, y\rangle = \mathrm{Cov}(x, y)$, and the ratio $\dfrac{\langle x, y\rangle}{\|x\|\|y\|}$ coincides with the *correlation coefficient* between x and y.

Projection onto a Subspace

The projection result established in Property 1 extends naturally to higher-dimensional subspaces. Let $M \subseteq L$ be a subspace with linearly independent basis $\{x_1, x_2, \dots, x_N\}$, and collect the basis elements into the random vector $\mathbf{x} = (x_1, \dots, x_N)'$. We say that $y_M \in M$ is the *projection* of $y \in L$ onto M if it minimizes $\|y - z\|$ over all $z \in M$.

Since any $z \in M$ takes the form $\mathbf{a}'\mathbf{x}$ for some $\mathbf{a} \in \mathbb{R}^N$, the projection solves

$$\min_{\mathbf{a}\in\mathbb{R}^N} \|y - \mathbf{a}'\mathbf{x}\|^2 = \min_{\mathbf{a}\in\mathbb{R}^N} \mathsf{E}\left[(y - \mathbf{a}'\mathbf{x})^2\right].$$

The first-order condition with respect to $\mathbf{a}$ yields

$$\mathsf{E}\left[(y - \mathbf{a}'\mathbf{x})\mathbf{x}\right] = \mathsf{E}(y\mathbf{x}) - \mathsf{E}(\mathbf{x}\mathbf{x}')\mathbf{a} = \mathbf{0}.$$

The Gram matrix $G = \mathsf{E}(\mathbf{x}\mathbf{x}')$ has (i, j) entry $\mathsf{E}(x_i x_j)$, the second cross-moment of the i-th and j-th basis payoffs; it is positive definite whenever the basis payoffs are linearly independent. Provided G is invertible, the unique minimizer is $\mathbf{a} = G^{-1}\,\mathsf{E}(y\mathbf{x})$, so

$$y_M = \mathbf{a}'\mathbf{x} = \mathsf{E}(y\mathbf{x})'G^{-1}\mathbf{x}.$$

The residual $z = y - y_M$ satisfies $\langle z, x_i \rangle = 0$ for every $i = 1, \ldots, N$. By linearity, this extends to $\langle z, x \rangle = 0$ for every $x \in M$, and the Pythagorean identity becomes $\|y\|^2 = \|y_M\|^2 + \|z\|^2$.

Property 3 (Projection onto a Subspace). *Let $M \subseteq L$ be a subspace with linearly independent basis $\{x_1, \ldots, x_N\}$, let $\mathbf{x} = (x_1, \ldots, x_N)'$, and assume the Gram matrix $G = \mathrm{E}(\mathbf{x}\mathbf{x}')$ is invertible. For any $y \in L$, the projection onto M is*

$$y_M = \mathrm{E}(y\mathbf{x})' G^{-1} \mathbf{x}.$$

The residual $z = y - y_M$ satisfies $\langle z, x \rangle = 0$ for all $x \in M$, and $\|y\|^2 = \|y_M\|^2 + \|z\|^2$.

The set of all payoffs in L that are orthogonal to M is called the *orthogonal complement* of M:

$$M^\perp = \{z \in L : \langle z, x \rangle = 0 \text{ for all } x \in M\}.$$

Since $M^\perp$ is closed under addition and scalar multiplication, it is itself a subspace of L. The projection theorem implies that every $y \in L$ admits the unique decomposition $y = y_M + z$ with $y_M \in M$ and $z \in M^\perp$, so $M \cap M^\perp = \{0\}$ and

$$L = M \oplus M^\perp.$$

This *orthogonal direct sum* decomposition is the key geometric fact underlying the structure of stochastic discount factors: as we will see in the next notebook, any valid SDF m decomposes as $m = x^* + e$, where $x^* \in X$ is the projection of m onto the subspace of traded payoffs and $e \in X^\perp$ is the orthogonal residual.

Linear Functionals

A central object in asset pricing is a *pricing functional* that assigns a price to every traded payoff. Under the law of one price, such a functional must be linear: the price of a portfolio equals the sum of the prices of its components. Understanding when and how linear functionals can be represented will therefore be essential for characterizing stochastic discount factors.

Given $x, y \in L$ and $\alpha, \beta \in \mathbb{R}$, a linear functional $f : L \to \mathbb{R}$ satisfies

$$f(\alpha x + \beta y) = \alpha f(x) + \beta f(y).$$

We say that the linear functional $f : L \to \mathbb{R}$ is bounded if

$$|f(x)| \leq M\|x\|$$

for all $x \in L$. In other words, the absolute value of the functional cannot grow infinitely for a finite x. A bounded linear functional is also called a continuous linear functional. The smallest M for which this inequality remains true is called the norm of f, i.e.,

$$\|f\| = \inf\{M : |f(x)| \leq M\|x\|, \text{ for all } x \in L\}.$$

For a given $m \in L$ and any $x \in L$, the functional

$$f(x) = \langle m, x \rangle = \mathrm{E}(mx) = \sum_{s=1}^{S} q(s)m(s)x(s)$$

is linear since

$$f(\alpha x + \beta y) = \mathrm{E}(m(\alpha x + \beta y)) = \alpha\, \mathrm{E}(mx) + \beta\, \mathrm{E}(my) = \alpha f(x) + \beta f(y).$$

Furthermore, the Cauchy-Schwartz inequality implies that

$$|f(x)| = |\langle m, x \rangle| \leq \|m\|\|x\|,$$

showing that the linear functional f is bounded and hence continuous. Since the previous inequality is an equality whenever x is proportional to m, we have that $\|m\|$ is the smallest bound of f, showing that $\|f\| = \|m\|$.

Conversely, consider a linear functional $f : L \to \mathbb{R}$. Its kernel $K = \{x \in L : f(x) = 0\}$ is a subspace of codimension one, so $K^{\perp}$ is one-dimensional: it is spanned by some nonzero z satisfying $\langle x, z \rangle = 0$ for all $x \in K$. Without loss of generality, assume that z has been appropriately scaled so that $f(z) = 1$.

Given any $x \in L$, we have that $x - f(x)z \in K$ since $f(x - f(x)z) = f(x) - f(x)f(z) = 0$. Moreover, $z \perp K$, implying that

$$0 = \langle x - f(x)z, z \rangle = \langle x, z \rangle - f(x)\langle z, z \rangle.$$

The previous expression implies that

$$f(x) = \frac{\langle x, z \rangle}{\langle z, z \rangle} = \langle x, m \rangle,$$

where $m = \dfrac{z}{\|z\|^2}$. The previous analysis is an important result known as the *Riesz representation theorem*.

Property 4 (Riesz Representation Theorem). *If $f : L \to \mathbb{R}$ is a bounded linear functional, there exists a unique vector $m \in L$ such that for all $x \in L$, $f(x) = \langle m, x \rangle$. Furthermore, we have $\|f\| = \|m\|$ and every m determines a unique bounded linear functional.*

The Riesz representation theorem has a direct and important consequence for asset pricing. Suppose that $p : L \to \mathbb{R}$ is a linear pricing functional satisfying the law of one price. Then the theorem guarantees the existence of a unique $m \in L$ such that

$$p(x) = \langle m, x \rangle = \mathrm{E}(mx)$$

for every payoff $x \in L$. The random variable m is called a *stochastic discount factor*, and characterizing its properties—positivity, uniqueness, and variance bounds—is the central task of asset pricing theory.