The Fundamental Theorem of Asset Pricing in L^p

Lorenzo Naranjo

March 2026

In the stochastic discount factor notebook we proved the equivalence PNA $\Leftrightarrow \exists\, m > 0$ in a finite state space using a direct separation argument in $\mathbb{R}^S$. That proof exploits the fact that the positive orthant $\mathbb{R}^S_{++}$ is open, which allows a global application of the Separating Hyperplane Theorem.

In an infinite-dimensional L^p space this shortcut fails: the cone of non-negative random variables has empty interior, so global separation is unavailable. This notebook presents the correct proof for the L^p setting, following the argument of Kabanov and Stricker (2001) and Clark (1993), both of which build on the classical results of Kreps (1981) and Yan (1980). The companion notebook FTAP with d Risky Assets presents the simpler case where the strategy space is finite-dimensional ($\theta \in \mathbb{R}^d$), in which simple no-arbitrage suffices.

The L^p Payoff Space

Let $p \in [1, \infty)$ and let q denote the conjugate exponent defined by $1/p + 1/q = 1$ (with $q = \infty$ when $p = 1$). Let $(\Omega, \mathcal{F}, P)$ be a probability space and let $L^p = L^p(\Omega, \mathcal{F}, P)$ denote the space of p-integrable random variables, equipped with the norm $\|x\|_p = \mathrm{E}(|x|^p)^{1/p}$. By the Riesz representation theorem, the dual space is $(L^p)^* = L^q$. The positive cone is

$$L^p_+ = \{x \in L^p : x \geq 0 \text{ a.s.}\},$$

and we write $x > 0$ to mean $x \geq 0$ a.s. and $P(x > 0) = 1$, i.e., x is strictly positive almost surely. We also write

$$L^p_{++} = \{x \in L^p : x > 0 \text{ a.s.}\}.$$

Unlike in $\mathbb{R}^S$, the positive cone L^p_+ has **empty interior**: every open ball around a non-negative function contains functions that dip below zero on sets of positive measure. For example, if

$x \in L^p_+$ and A has positive probability, then $x - \varepsilon \mathbf{1}_A$ is negative on A for large enough $\varepsilon > 0$. This single fact is the source of all the additional difficulty in the infinite-dimensional proof.

> **Note**
>
> When $p = 2$, L^2 is a Hilbert space under the inner product $\langle x, y \rangle = \mathsf{E}(xy)$, with $(L^2)^* = L^2$ (self-duality). For $p \neq 2$, L^p is a Banach space with $(L^p)^* = L^q$, $q \neq p$. The proof below applies uniformly to all $p \in [1, \infty)$, though certain steps simplify for $p = 2$.

The Trading Model

Let $X \subseteq L^p$ be a closed subspace of traded payoffs, and let $\pi : X \to \mathbb{R}$ be a continuous linear pricing functional (the law of one price). Define the *cone of claims attainable at non-positive cost* by

$$K = \{y \in L^p : \exists\, x \in X \text{ with } \pi(x) \leq 0 \text{ and } y \leq x \text{ a.s.}\}.$$

Economically, K is the set of payoffs that can be super-replicated with initial cost at most zero. Any payoff that is dominated by a zero-cost traded payoff belongs to K. In particular, $-L^p_+ \subseteq K$: every non-positive payoff is in K (take $x = 0$). A *free lunch* is a sequence $(y_n) \subseteq K$ that converges in L^p to some $y \in L^p_+$ with $y \neq 0$: the payoffs y_n approximate an arbitrage in the L^p sense, even if no exact arbitrage exists. Denote by $\bar{K}$ the closure of K in L^p; then $\bar{K}$ is a closed convex cone satisfying $-L^p_+ \subseteq \bar{K}$.

Assumption 1 (No Free Lunch (L^p))**.**

$$\bar{K} \cap L^p_+ = \{0\}.$$

This says no sequence of dominated zero-cost payoffs can converge (in L^p) to a non-negative, non-zero payoff. The weaker condition $K \cap L^p_+ = \{0\}$ (no exact free lunch in K) is *not* sufficient in infinite dimensions: one can construct examples of markets where no exact arbitrage in traded assets exists yet there are free lunches — sequences of strategies whose downside vanishes while their upside remains bounded. The No Free Lunch condition rules out both.

Note

Clark (1993) formulates a slightly weaker no-free-lunch axiom (his A.5): only strictly positive payoffs ($x > 0$ a.s.) are required to lie outside $\bar{K}$, i.e., $\bar{K} \cap L^p_{++} = \emptyset$. The condition above is stronger: it rules out all non-zero non-negative payoffs. The stronger form simplifies Step 2 of the proof, which takes $x = \mathbf{1}_A$ for sets A of positive probability; since $\mathbf{1}_A$ is not strictly positive when $P(A^c) > 0$, Clark's A.5 alone does not guarantee $\mathbf{1}_A \notin \bar{K}$. Clark resolves this via a supporting set argument for each $x \in L^p_{++}$; the condition above instead allows Step 2 to proceed directly.

The Fundamental Theorem

Property 1 (Fundamental Theorem of Asset Pricing). *Assume:*

1. *$X \subseteq L^p$ is a closed subspace of traded payoffs,*
2. *$\pi : X \to \mathbb{R}$ is a continuous linear pricing functional (law of one price),*
3. *no free lunch holds: $\bar{K} \cap L^p_+ = \{0\}$,*
4. *there exists a traded numeraire $x_0 \in X$ with $x_0 > 0$ a.s. and $\pi(x_0) > 0$.*

Then there exists $m \in L^q$ with $m > 0$ a.s. such that

$$\pi(x) = \mathrm{E}(mx) \quad \text{for all } x \in X.$$

The random variable m is a strictly positive stochastic discount factor. The measure $\tilde{P}$ defined by $d\tilde{P}/dP = m/\mathrm{E}(m)$ satisfies $\tilde{P} \sim P$ (the two measures share the same null sets). In this static setting, $\tilde{P}$ is an equivalent pricing measure; in dynamic models it plays the role of an equivalent martingale measure. The SDF lies in the dual space L^q: when $p = 1$, $m \in L^\infty$ is bounded; when $p = 2$, $m \in L^2$ is square-integrable.

Proof Tools

The proof uses two results. The first holds in any L^p space; the second is a measure-theoretic fact about families of measures.

Roadmap: Hahn-Banach gives point-wise separators in $L^q = (L^p)^*$; Halmos-Savage reduces the resulting uncountable family to a countable subfamily that preserves null sets.

Lemma 0.1 (Hahn-Banach Separation in L^p). *Let $C \subseteq L^p$ be a non-empty closed convex set, and let $x \notin C$. Then there exists $z \in L^q$ such that*

$$\mathsf{E}(z\, y) \leq 0 < \mathsf{E}(z\, x) \quad \textit{for all } y \in C,$$

provided C is a cone containing 0.

This follows from the Hahn-Banach separation theorem applied in the locally convex space L^p, together with the duality $(L^p)^* = L^q$. The separator is a bounded linear functional on L^p; by the Riesz representation theorem it is represented as $f(\cdot) = \mathsf{E}(z\, \cdot)$ for a unique $z \in L^q$. When $p = 2$, the self-duality $(L^2)^* = L^2$ means the separator is already an element $z \in L^2$: the Hilbert-space projection theorem and Hahn-Banach collapse into one step, and no separate identification via Riesz is needed. This is the key advantage of the L^2 setting: the dual is the same space, whereas for $p \neq 2$ the separator lives in a different space $L^q \neq L^p$.

Lemma 0.2 (Halmos-Savage Theorem). *Let $\{\mu_\alpha : \alpha \in I\}$ be a family of measures on $(\Omega, \mathcal{F})$, each absolutely continuous with respect to P. If for every measurable A with $P(A) > 0$ there exists some α with $\mu_\alpha(A) > 0$, then there exists a countable subfamily $\{\mu_{\alpha_n}\}_{n\geq 1}$ with the same null sets as the full family.*

This theorem reduces an uncountable family of measures to a countable one without losing information about null sets. It is the step that has no finite-dimensional analogue: in $\mathbb{R}^S$ a single global separator suffices, but in L^p we can only separate point-by-point and must then stitch the resulting separators together.

Proof of Property 1

The proof has four steps. Step 1 applies Hahn-Banach to each indicator function $\mathbf{1}_A$, building a family of non-negative separators indexed by positive-measure sets that collectively charges every non-null event. Step 2 uses Halmos-Savage to extract a countable equivalent subfamily. Step 3 averages that subfamily with geometrically-decaying weights to form a strictly positive density $\rho \in L^q$, normalized to m. Step 4 verifies that m prices correctly.

Step 1: Build a covering family of separators. For each measurable A with $P(A) > 0$, apply Lemma 0.1 to $\mathbf{1}_A \in L_+^p$: since $\mathbf{1}_A \notin \bar{K}$ by no free lunch, there exists $z_A \in L^q$ with

$$\mathsf{E}(z_A\, y) \leq 0 \quad \text{for all } y \in \bar{K}, \qquad \mathsf{E}(z_A\, \mathbf{1}_A) > 0.$$

Since $-L_+^p \subseteq \bar{K}$, for any $w \geq 0$ we have $\mathsf{E}(z_A\, w) \geq 0$. Taking $w = \mathbf{1}_{\{z_A<0\}}$ gives $\mathsf{E}(z_A\, \mathbf{1}_{\{z_A<0\}}) \geq 0$, but $z_A < 0$ on $\{z_A < 0\}$, so $P(z_A < 0) = 0$, i.e., $z_A \geq 0$ a.s. Normalize to $\|z_A\|_q = 1$ (when $q < \infty$) or $\|z_A\|_\infty = 1$ (when $q = \infty$). Define the measure $\mu_A(B) := \mathsf{E}(z_A\, \mathbf{1}_B)$. Then $\mu_A(A) = \mathsf{E}(z_A\, \mathbf{1}_A) > 0$, so the family $\{\mu_A : P(A) > 0\}$ charges every positive-P-measure set by construction.

Step 2: Extract a countable equivalent subfamily. Each individual separator z_A may vanish on most of Ω; what we need is a *single* density that is positive everywhere. By Lemma 0.2, there exists a countable sequence of sets $\{A_n\}_{n\geq 1}$ (with $P(A_n) > 0$ for each n) such that the measures $\{\mu_{A_n}\}$ have the same null sets as the full family $\{\mu_A\}$. In particular, if $P(B) > 0$ then $\mathsf{E}(z_{A_n}\, \mathbf{1}_B) > 0$ for some n.

Step 3: Construct the SDF. We average the countable family with geometrically-decaying weights to produce a single density. Define

$$\rho = \sum_{n=1}^{\infty} 2^{-n} z_{A_n}.$$

When $q < \infty$: the triangle inequality gives $\|\rho\|_q \leq \sum_{n=1}^{\infty} 2^{-n} \|z_{A_n}\|_q = 1$, so $\rho \in L^q$. When $q = \infty$ (i.e., $p = 1$): since $0 \leq z_{A_n} \leq 1$ a.s., we have $0 \leq \rho \leq 1$ a.s. By Step 2, $\rho > 0$ a.s.: if $P(\rho = 0) > 0$ then there is a set B with $P(B) > 0$ where every summand $2^{-n} z_{A_n}$ vanishes, so $\mathsf{E}(z_{A_n}\, \mathbf{1}_B) = 0$ for all n, contradicting the null-set equivalence. Since $\rho \in L^q$ and P is a probability measure, $0 < \mathrm{E}(\rho) \le \|\rho\|_q \le 1$. Set $m = \rho / \mathrm{E}(\rho) \in L^q$; then $m > 0$ a.s. and $\mathrm{E}(m) = 1$.

Step 4: m prices all assets correctly. Define $\ell : X \to \mathbb{R}$ by $\ell(x) := \mathrm{E}(mx)$, well-defined by Hölder's inequality.

Claim 1: $\ell(y) \le 0$ for all $y \in \bar{K}$. Each z_{A_n} satisfies $\mathrm{E}(z_{A_n}\, y) \le 0$ for all $y \in \bar{K}$. By Hölder, $|\,\mathrm{E}(z_{A_n}\, y)| \le \|z_{A_n}\|_q \|y\|_p = \|y\|_p$, so the series $\sum_{n=1}^{\infty} 2^{-n}\, \mathrm{E}(z_{A_n}\, y)$ is dominated by $\|y\|_p$ and converges. Swapping sum and expectation by dominated convergence:

$$\mathrm{E}(\rho\, y) = \sum_{n=1}^{\infty} 2^{-n}\, \mathrm{E}(z_{A_n}\, y) \le 0.$$

Dividing by $\mathrm{E}(\rho) > 0$ gives $\ell(y) = \mathrm{E}(my) \le 0$ for all $y \in \bar{K}$.

Proportionality. For any $x \in \ker \pi$, both x and $-x$ lie in $K \subseteq \bar{K}$, so Claim 1 gives $\ell(x) \le 0$ and $\ell(-x) \le 0$, hence $\ell(x) = 0$. Thus $\ker \pi \subseteq \ker \ell$. Since $\pi \not\equiv 0$ (assumption 4 gives $\pi(x_0) > 0$), the quotient $X / \ker \pi \cong \mathbb{R}$ is one-dimensional, and any functional vanishing on $\ker \pi$ is a scalar multiple of π. Hence $\ell = c\pi$ for some $c \in \mathbb{R}$.

$c > 0$. The numeraire x_0 satisfies $x_0 > 0$ a.s. and $m > 0$ a.s., so $\ell(x_0) = \mathrm{E}(mx_0) > 0$. Therefore $c = \ell(x_0)/\pi(x_0) > 0$.

Rescale. Set $\hat{m} = m/c$. Then $\hat{m} \in L^q$, $\hat{m} > 0$ a.s. and

$$\mathrm{E}(\hat{m}\, x) = \frac{\ell(x)}{c} = \frac{c\,\pi(x)}{c} = \pi(x) \qquad \text{for all } x \in X.$$

□

Why the Finite-Dimensional Proof Is Simpler

The proof above has three steps that vanish entirely in $\mathbb{R}^S$:

No need for point-wise separation. In $\mathbb{R}^S$, the positive orthant $\mathbb{R}^S_{++}$ is open, so the relevant pricing set and $\mathbb{R}^S_{++}$ are disjoint convex sets with one of them open. The Separating Hyperplane Theorem applies *globally* and produces a single separator ϕ in one shot. In L^p, the positive cone has empty interior, so global separation is not available and one must separate each point $x \in L^p_+$ individually.

No need for Halmos-Savage. The global separator in $\mathbb{R}^S$ is already a single vector $\phi \in \mathbb{R}^S$. In L^p, the point-wise separators $\{z_x\}$ form an uncountable family, and Halmos-Savage is the mechanism that stitches them into a single density ρ.

Simple no-arbitrage is not enough. In $\mathbb{R}^S$, every linear subspace is automatically closed, so K is automatically closed and no-arbitrage ($K \cap \mathbb{R}^S_+ = \{0\}$) suffices. In L^p, the cone K need not be closed: one can construct markets with no exact arbitrage yet with *approximate* arbitrages — sequences of dominated zero-cost payoffs converging to a non-negative, non-zero payoff (a *free lunch*). The correct infinite-dimensional hypothesis is therefore the stronger *no free lunch* condition $\bar{K} \cap L^p_+ = \{0\}$.

In this L^p model, the **no free lunch** (NFL) condition is genuinely stronger than simple **no arbitrage** (NA). The precise difference stems from the possible non-closedness of K. Simple no-arbitrage requires only

$$K \cap L^p_+ = \{0\},$$

i.e., no exact non-negative non-zero payoff is attainable at non-positive cost. No free lunch imposes the strictly stronger requirement

$$\bar{K} \cap L^p_+ = \{0\},$$

where $\bar{K}$ denotes the closure of K in the L^p topology. When K is not closed, it is possible to satisfy NA while violating NFL: there exist sequences of strategies in K that converge in L^p to a strictly positive payoff, even though no single strategy in K delivers a non-negative non-zero payoff exactly.

The companion notebook FTAP with d Risky Assets presents the simpler proof for finite-dimensional strategy spaces. That notebook also explains why the two settings differ: the finite dimensionality of $\mathbb{R}^d$ supplies compactness (via Bolzano-Weierstrass) that forces NA to imply closedness, and the L^1/L^∞ duality (when $p = 1$) yields a bounded density $d\tilde{P}/dP \in L^\infty$, whereas the present L^p proof gives only $m \in L^q$.

References

Clark, Stephen A. 1993. "The Valuation Problem in Arbitrage Price Theory." *Journal of Mathematical Economics* 22 (5): 463–78.

Kabanov, Yuri M., and Christophe Stricker. 2001. "A Teachers' Note on No-Arbitrage Criteria." In *Séminaire de Probabilités XXXV*, vol. 1755. Lecture Notes in Mathematics. Springer.

Kreps, David M. 1981. "Arbitrage and Equilibrium in Economies with Infinitely Many Commodities." *Journal of Mathematical Economics* 8 (1): 15–35.

Yan, Jia-An. 1980. "Caractérisation d'une Classe d'ensembles Convexes de L^1 Et H^1." In *Séminaire de Probabilités XIV*, vol. 784. Lecture Notes in Mathematics. Springer.