

A Nonlinear Spectral Fundamental Theorem of Asset Pricing

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

Reader-Friendly Subtitle

Extending no-arbitrage theory beyond linear pricing rules.

Technical Strapline

No-free-lunch is equivalent to existence of a consistent nonlinear spectral pricing kernel.

Executive Summary (Non-Technical)

Classical no-arbitrage theory is one of finance's strongest foundations, but much of it is built around linear structures. Real markets frequently behave in ways that are effectively nonlinear.

This paper proposes a spectral extension of FTAP that keeps the core no-arbitrage logic while allowing mode-dependent nonlinear pricing effects. The goal is to preserve rigor while expanding modeling reach.

If successful, the result provides a unified bridge between pricing and allocation: both become projections of the same spectral object rather than separate frameworks.

This paper does not claim to replace all linear models. It provides a strict extension where the linear case appears as a special regime.

Abstract

Classical FTAP characterizes no-arbitrage in linear pricing systems. We develop a nonlinear spectral extension in which state prices are represented as mode-dependent functionals over a spectral decomposition of market states. We prove equivalence between no-free-lunch conditions and existence of a consistent nonlinear spectral pricing kernel, yielding a unified theorem that nests CAPM-like linear cases as a strict special regime.

1. Problem

The classical FTAP (Harrison-Kreps 1979, Delbaen-Schachermayer 1994) characterizes no-arbitrage through existence of an equivalent martingale measure (EMM). This is powerful but fundamentally linear: pricing is a linear functional on payoffs.

Real markets exhibit state-dependent liquidity, nonlinear transaction costs, and mode-dependent risk premia that break the linear pricing assumption. The problem is: can we extend the FTAP equivalence to a nonlinear spectral setting where each mode k has its own pricing rule $\Pi_k(\cdot)$, while preserving the core no-arbitrage pricing kernel duality?

2. Setup

2.1 Market State and Mode Decomposition

Let $X = (X_1, \dots, X_n)$ be asset returns. Decompose into orthogonal spectral modes: $c_k = v_k^\top X$ where v_k are eigenvectors of the return covariance Σ .

Each mode c_k has variance λ_k and risk premium $\pi_k = \mathbb{E}[c_k]$.

2.2 Nonlinear Spectral Pricing Functional

Definition 1 (Spectral Pricing Functional). A nonlinear spectral pricing system is a collection $\{\Pi_k : L^2(\Omega) \rightarrow \mathbb{R}\}_{k=1}^K$ such that: - (P1) Π_k is monotone: $Y \geq Z$ a.s. $\Rightarrow \Pi_k(Y) \geq \Pi_k(Z)$ - (P2) $\Pi_k(1) = 1$ (normalization) - (P3) Π_k is convex (sublinear pricing admits transaction costs)

The aggregate price of a payoff Y is:

$$\Pi(Y) = \sum_{k=1}^K w_k \Pi_k(v_k^\top Y)$$

2.3 No-Free-Lunch in Mode Space

Definition 2 (Spectral No-Free-Lunch, SNFL). There is no portfolio θ such that $\Pi(\theta^\top X) < 0$ and $\theta^\top X \geq 0$ a.s. with $P(\theta^\top X > 0) > 0$.

Connection to existing kernel: the linear case $\Pi_k(c_k) = \lambda_k^{-1} \pi_k \cdot c_k$ recovers the Pricing Is Allocation identity $\pi_k = \lambda w_k^{\text{mkt}} v_k$ from PricingAllocation/ModeRiskPremium.lean.

3. Main Theorem

Theorem Candidate 1 (Nonlinear Spectral FTAP). The following are equivalent: 1. SNFL holds. 2. There exists a consistent nonlinear spectral pricing kernel $\{\Pi_k\}$ satisfying (P1)-(P3) such that $\Pi(\theta^\top X) \geq 0$ for all admissible θ .

Corollary (Linear Recovery). If each Π_k is linear, SNFL reduces to classical NFLVR and $\Pi_k(c_k) = \mathbb{E}^Q[c_k]$ for some EMM Q .

Theorem Candidate 2 (Mode-Wise Decomposition). Under SNFL, the per-mode pricing rule satisfies:

$$\Pi_k(c_k) = \mathbb{E}[c_k] + \phi_k(\text{Var}(c_k), \kappa_k, \dots)$$

where ϕ_k is a convex nonlinear adjustment depending on mode-specific tail statistics.

4. Proof Sketch

1. **Dual construction.** Work in $L^2(\Omega)$ mode-decomposed. The SNFL condition defines a convex cone of nonnegative payoffs. By Hahn-Banach, a separating functional exists.
2. **Mode-wise consistency.** The separating functional decomposes into per-mode components Π_k by orthogonality of modes. Monotonicity (P1) follows from the cone structure. Convexity (P3) follows from the sublinearity of the separating functional.
3. **Normalization and uniqueness.** Normalization (P2) is fixed by the risk-free rate. Uniqueness fails in general (incomplete markets), but a minimal pricing system exists.
4. **Linear reduction.** If Π_k is additive, Riesz representation gives $\Pi_k(Y) = \mathbb{E}^Q[Y]$ for some measure Q , recovering classical EMM.

5. Empirics/Simulation

5.1 Toy Nonlinear Market

- Two-asset market with state-dependent transaction costs.
- Compare: (a) linear EMM pricing, (b) nonlinear spectral pricing with mode-specific convexity.
- Measure arbitrage residuals and pricing consistency under stress.

5.2 Empirical Equity Modes

- S&P 500 PCA modes with historical data.
- Estimate mode-specific nonlinear pricing adjustments from option-implied information.
- Test whether nonlinear corrections improve hedging performance.

6. Limits

- **Non-uniqueness:** in incomplete markets, the nonlinear kernel is not unique; only bounds are available.
- **Computational cost:** nonlinear pricing requires mode-by-mode optimization rather than a single EMM.
- **Decomposition dependence:** results depend on the quality and stability of the spectral decomposition.
- **Dynamic extension:** the static version is proved here; dynamic (multi-period) extension is future work.

7. Related Work

- **Classical FTAP:** Harrison-Kreps (1979), Delbaen-Schachermayer (1994, 1998).
- **Convex risk/pricing duality:** Föllmer-Schied (2016), Frittelli-Rosazza Gianin (2002).
- **Nonlinear pricing:** El Karoui-Quenez (1995), Peng (2004) on g-expectations.
- **Spectral pricing:** our Pricing Is Allocation paper — the linear spectral kernel is the base case.
- **Mode risk premia:** our Spectral Trading paper — mode Sharpe ratios and no-free-lunch per mode.

8. Cross-Paper Connections

- **Topological Obstructions (paper 9):** nonlinear pricing forms may have cohomological obstructions that prevent global consistency. The SNFL condition here is a local statement; paper 9 shows when local SNFL cannot extend globally on nontrivial state manifolds.
- **Information Geometry Bridge (paper 6):** the metric compatibility theorem provides bounds on how nonlinear pricing distortions affect the Fisher geometry of the underlying model family. The bridge map Φ interleaves with Π_k .
- **Decision Functional Approximation (paper 8):** the nonlinear pricing functional Π is itself a decision functional. Paper 8 provides approximation rates for Π when evaluated in spectral coordinates with finite mode budget.
- **Minimal Sufficient State (paper 4):** the minimal spectral budget for preserving nonlinear pricing consistency is a special case of the K^* result from paper 4 applied to the functional Π .

9. Concrete Example: Two-Asset Nonlinear Market

9.1 Setup

Two assets with returns X_1, X_2 . Covariance Σ has eigenvectors v_1, v_2 with eigenvalues $\lambda_1 > \lambda_2$. Spectral modes: $c_1 = v_1^\top X$, $c_2 = v_2^\top X$.

Linear regime (classical FTAP): $\Pi_k^{\text{lin}}(c_k) = \mathbb{E}^Q[c_k]$ for EMM Q . Pricing is additive: $\Pi^{\text{lin}}(Y) = w_1 \mathbb{E}^Q[c_1 \cdot v_1^\top Y] + w_2 \mathbb{E}^Q[c_2 \cdot v_2^\top Y]$.

Nonlinear extension (with transaction costs): Suppose mode 1 (dominant variance) faces proportional transaction cost κ_1 , while mode 2 (low variance) is frictionless. Define:

$$\begin{aligned}\Pi_1(c_1) &= \mathbb{E}^Q[c_1] + \kappa_1 \mathbb{E}^Q[|c_1|] \\ \Pi_2(c_2) &= \mathbb{E}^Q[c_2]\end{aligned}$$

Π_1 is convex, monotone, and normalized. The aggregate price is:

$$\Pi(Y) = w_1 \Pi_1(v_1^\top Y) + w_2 \Pi_2(v_2^\top Y)$$

9.2 SNFL Verification

For any portfolio θ , decompose $\theta^\top X = a_1 c_1 + a_2 c_2$.

$$\Pi(\theta^\top X) = w_1 [\mathbb{E}^Q[a_1 c_1] + \kappa_1 |a_1| \mathbb{E}^Q[|c_1|]] + w_2 \mathbb{E}^Q[a_2 c_2].$$

If $\theta^\top X \geq 0$ a.s. with positive probability of strict positivity, then by classical no-arbitrage on Q , $\mathbb{E}^Q[\theta^\top X] > 0$. The nonlinear correction is nonnegative, so $\Pi(\theta^\top X) > 0$. SNFL holds.

9.3 Hedging Implications

The nonlinear pricing functional produces a bid-ask spread in mode space:

- **Bid price:** $\Pi_1^{\text{bid}}(c_1) = \mathbb{E}^Q[c_1] - \kappa_1 \mathbb{E}^Q[|c_1|]$
- **Ask price:** $\Pi_1^{\text{ask}}(c_1) = \mathbb{E}^Q[c_1] + \kappa_1 \mathbb{E}^Q[|c_1|]$

The mode-specific spread $2\kappa_1 \mathbb{E}^Q[|c_1|]$ scales with mode volatility. Higher-variance modes get wider spreads — a structural prediction testable on real equity factor data.

Corollary (Mode-Selective Hedging). The cheapest hedge of a payoff Y under nonlinear pricing uses only frictionless modes when possible:

$$\theta^* = \arg \min_{\theta} \Pi(\theta^\top X - Y) \quad \text{subject to} \quad \theta^\top X \geq Y \text{ a.s.}$$

This optimization naturally routes hedging activity away from mode 1 toward mode 2, consistent with the empirical observation that hedging concentrates in liquid factors.

10. Outlook

- **XVA integration:** nonlinear mode-dependent pricing kernels can encode CVA/DVA/FVA adjustments per spectral mode rather than as portfolio-level add-ons.
- **Incomplete-market bounds:** the nonlinear FTAP naturally yields super/sub-hedging price bounds. The two-asset example demonstrates that bid-ask spreads decompose mode-wise.
- **Dynamic extension:** extend to multi-period setting with mode-dependent filtration.
- **Lean formalization:** the Hahn-Banach separation step and mode-decomposition are natural targets for LeanProofs/NonlinearFTAP/. The two-asset SNFL verification (Section 9.2) is an explicit first formalization target.
- **Empirical priority:** test the mode-selective hedging prediction on real equity PCA factors with transaction cost estimates from tick data.