Projective Kernel Bands and Gain–Loss Bounds in Incomplete Markets

Dimitris Margaritis · Maryam Hasannasab

University of Auckland Business School

EWEPA 2026 · This version: June 2026

Questions

- What does the math used to measure the efficiency of a farm or a manufacturing plant have to do with how Wall Street hedge funds price complex multi-million dollar financial derivatives?
- Why would an investor ever intentionally buy a portfolio that wastes money or gives them more output than they actually need?
- What happens when the financial system doesn't offer a perfectly tailored, distinct asset for every single possible future state of the world?

Incomplete Markets: Wide Price Intervals

No-arbitrage prices a non-spanned payoff only to an interval — often too wide to inform.

Two ways to restrict admissible pricing kernels:

- **Gain–loss bounds** (Bernardo–Ledoit 2000): cap the gain–loss ratio ⇔ bound the statewise ratio m_s / m^*_s of a candidate kernel to a benchmark kernel, aka a rational investor would not insist on an impossible good deal.
- **Directional distance** (Chambers–Färe 2011): instead of using subjective risk preferences, they leverage the DDF treating payoff generation as a "production technology" to bound the prices of assets that lie outside the span of the current market.
- While CF determine the maximum and minimum prices before an asset becomes a "free lunch" (arbitrage), BL recognize that pure arbitrage is rare, so they bound prices by ruling out investment opportunities that are "too good to be true."
- If you imagine pricing bounds as a fence holding back unrealistic prices, CF replace the fence at the absolute outer edge of what the physical asset structures in the market can replicate (pure arbitrage). BL pull the fence inward, arguing that a price is economically invalid long before it becomes a literal arbitrage opportunity if it provides an absurdly high return for almost no downside.

This paper: a finite-state bridge.

- The probability-weighted inner product $\langle m, z \rangle = \Sigma\, p_s\, m_s\, z_s$ pairs payoff vectors and pricing kernels in the same finite-dimensional space — letting benchmark rays and direction vectors be compared geometrically while keeping their distinct economic roles.

Incomplete Markets

- In a complete market, there's a unique asset payoff for every single future economic state, creating a single, unique pricing kernel.
- In reality market are incomplete, there are far more future states of the world than there are tradable assets which means most future risks cannot be perfectly hedged.
- Instead of one unique pricing kernel, there are infinitely many stochastic discount factors that can perfectly price the existing traded assets.
- In such a wide and ambiguous kernel space, you can pick extreme highly volatile kernels that price existing assets but assign absurd values to untraded assets.
- The BL benchmark solution, pick a macroeconomic reference ray m* and narrow down this infinite subspace by enforcing a strict Loss-Gain ratio bound across all finite states. This allows a tighter interval of fair prices for any newly introduced untraded asset.
- Similarly, in an incomplete market, the payoff space does not span the entire universe of possibilities. The CF solution is local and directional, use a payoff direction vector to cut straight through the empirical gaps of the market, evaluating efficiency relative only to what is tangible achievable. Rather than producing an interval, the CF dual profit function yields a unique upper- or lower-bound shadow price based on the cost of the closest portfolio on the market frontier.

The Probability weighted duality bridge map

- **Probability-Weighted Pairing**: The primal payoff space and dual kernel space pair under the probability-weighted inner product $\langle m, z \rangle = \Sigma\, p_s\, m_s\, z_s$ establishing the fundamental mathematical link between asset payoffs and market prices.
- **Primal Trajectory Selection**: The Chambers-Färe DDF acts as a support function in payoff space, where an analyst selects a direction vector to define the exact path and translation trajectory toward the market efficiency frontier.
- **Dual Domain Truncation**: The Bernardo-Ledoit gain-loss parameter establishes bounding hyperplanes that cut off extreme regions of the dual space, compressing the infinite martingale domain into a compact, convex pricing polytope.
- **Synthesis and Evaluation**: The probability weights bridge the economic frameworks, where the chosen primal direction sets the optimization slope and the dual parameter boxes in the feasible region, yielding a tight asset pricing corridor that prevents both structural inefficiency and approximate arbitrage.

Three Results

1

Exact recovery

"We didn't approximate BL—we recovered it."

When the reference ray is the Bernardo–Ledoit benchmark ray, b = m*, our band reproduces the finite-state BL admissible kernel set exactly — stronger than merely b = m*.

2

Chambers–Färe link

"No matter which admissible kernel you choose, replicated assets have one price."

On the marketed span every admissible kernel agrees, so bounds collapse to the unique no-arbitrage price; off-span the band gives explicit LP bounds.

3

Projective geometry

“"No matter which admissible kernel you choose, replicated assets have one price."

The band depends only on the positive ray [b] (scale-invariant). Bounds are linear programs with an interpretable dual; the impact of changing the benchmark ray is quantified by the Hilbert projective distance, a measure of differences in relative state valuation.

The Finite-State Economy

States & objects

- $\Omega = \{1, \ldots, S\}$; $p_s > 0$, $\Sigma\, p_s = 1$
- Payoff $z \in \mathbb{R}^S$; pricing kernel (SDF) $m \in \mathbb{R}^S$
- Value: $\langle m, z \rangle = \Sigma\, p_s\, m_s\, z_s$, $P = \mathrm{diag}(p_1, \ldots, p_S)$

Traded assets & pricing

- J assets — payoff matrix $A \in \mathbb{R}^{S \times J}$, prices $q \in \mathbb{R}^J$ $\Rightarrow$ $A'P\, m = q$ *(*Find all kernels that correctly reproduce observed market prices. Each assets creates one linear restriction).
- Riskless asset traded $\Rightarrow$ $\langle m, 1 \rangle = q_f = 1 / R_f$ (The traded risk-free asset anchors the overall level of the pricing kernel. While our projective band ultimately depends only on relative state weights, actual asset prices require a normalization, and the bond provides that normalization through its known price)

Admissible pricing kernels

- $M = \{ m \in \mathbb{R}^S_+ : A'P\, m = q \}$ is convex and compact
- Assumption 1 *(no arbitrage)*: M contains a strictly positive kernel $m \gg 0$
- Marketed payoffs priced uniquely; non-marketed payoffs give a wide interval $[\, \pi(z), \bar{\pi}(z) \,]$

The Projective Kernel Band

PROJECTIVE KERNEL BAND

$$M(L\,;\,b) = \{\, m \in M \;:\; \exists\, r > 0,\quad r\,b_s \le m_s \le L\,r\,b_s \;\; \text{for all } s \,\}$$

equivalently $\quad [\,\max_s (m_s/b_s)\,] \,/\, [\,\min_s (m_s/b_s)\,] \;\le\; L$

- Reference vector $b \in \mathbb{R}^S_+$, ceiling $L \in [1, \infty]$; $M(\infty\,;\,b) = M$
- **Projective:** $M(L\,;\,\lambda b) = M(L\,;\,b)$ for every $\lambda > 0$ — only the positive ray $[\,b\,]$ matters (rescaling b leaves the set unchanged, so only the positive ray through b matters, not its normalization)
- Bounds are linear programs:
- Upper bound: $\bar{\pi}(z\,;\,L,\,b) = \sup \langle m,\, z \rangle$, (Choose the kernel that makes the call as expensive as possible,)
- Lower bound: $\pi(z\,;\,L,\,b) = \inf \langle m,\, z \rangle$, (Choose the kernel that makes the call as cheap as possible,)
- Width $W = \bar{\pi} - \pi$

Properties of the Band

Prop. 1 — Compactness & existence

For finite $L > 1$, if $M(L\,;\,b) \neq \emptyset$ it is convex and compact — so upper and lower bounds are finite and attained for every payoff z.

Prop. 2 — Monotonicity in L

$1 \le L_1 \le L_2 \;\Rightarrow\; M(L_1\,;\,b) \subseteq M(L_2\,;\,b)$. Bands nest: the interval widens monotonically as the ceiling rises so L is a genuine tightness dial.

Prop. 3 — Full no-arbitrage limit

$M(\infty\,;\,b) = M$ for every b. The $L = \infty$ bounds are exactly the full no-arbitrage interval — independent of the reference ray.

Prop. 4 — Convergence

As $L \to \infty$, $\bar{\pi}(z\,;\,L,\,b) \to \bar{\pi}_{NA}(z)$ and $\pi(z\,;\,L,\,b) \to \pi_{NA}(z)$ for every reference and every payoff.

Exact Recovery of Bernardo–Ledoit

THEOREM 1 · BERNARDO–LEDOIT RECOVERY

$$b = m^* \quad \Rightarrow \quad M(L\,;\,b) = M_{BL}(L\,;\,m^*)$$

- BL benchmark — a strictly positive kernel $m^* \in \mathbb{R}^S_+$ with gain–loss set $M_{BL}(L\,;\,m^*) = \{\, m \in M : \exists\, 0 < \alpha \le \beta,\ \beta/\alpha \le L,\ \alpha\, m^*_s \le m_s \le \beta\, m^*_s \,\}$
- The set equality — hence both bounds $\bar{\pi}, \pi$ — holds for any reference on the BL ray $[\, m^* \,]$, for every payoff z.
- **Stronger than $b = m^*$:** projective invariance means any positive rescaling of the benchmark yields the same admissible set (Remark 1).
- We are not offering an alternative to Bernardo–Ledoit — we embed their admissible set as a special case and make its geometry, scale invariance, and computation transparent.

Connection to Chambers–Färe

Direction g — payoff space

- $D_C(z\,;\,g) = \sup\{\,\delta \in \mathbb{R} : z + \delta g \in C\,\}$
- By varying g, we are physically choosing a trajectory to push or expand the unpriced payoff toward the boundary of the existing production set.

Reference b — kernel space

- Restricts SDFs via the statewise spread of m/b
- Same $\mathbb{R}^S$ only because finite-state; economic roles differ.

On the marketed span $\mathcal{X}$: all admissible kernels agree $\Rightarrow$ $\bar{\pi} = \pi = \vartheta' q$ (unique no-arbitrage price)

- **Off-span:** $M(\infty\,;\,b) = M$ recovers the full no-arbitrage interval, while CF directional valuation stays direction-dependent.
- **Reference sensitivity (Thm 4):** $M(L\,;\,b) \subseteq M(\kappa(b,c)\,L\,;\,c)$, with $\log \kappa(b,c)$ the Hilbert projective distance — the reference ray is not innocuous (changing the reference ray costs a multiplicative factor equal to the exponential of the Hilbert projective distance).

Numerical Illustration: A European Call

$S_0 = 100$, $S_1 = (90, 100, 115)$, bond + stock span 2 of 3 states so the call is not replicated. Call K = 100 → z = (0, 0, 15).
Admissible reference b = (0.7050, 0.7061, 1.5282) prices both assets obtained as solution A'Pb=q; reference price $\langle b, z \rangle = \Sigma\, p_s\, b_s\, z_s = 6.88$ option value implied by the benchmark valuation rule b (when L=1 the projection band only contains the reference ray).

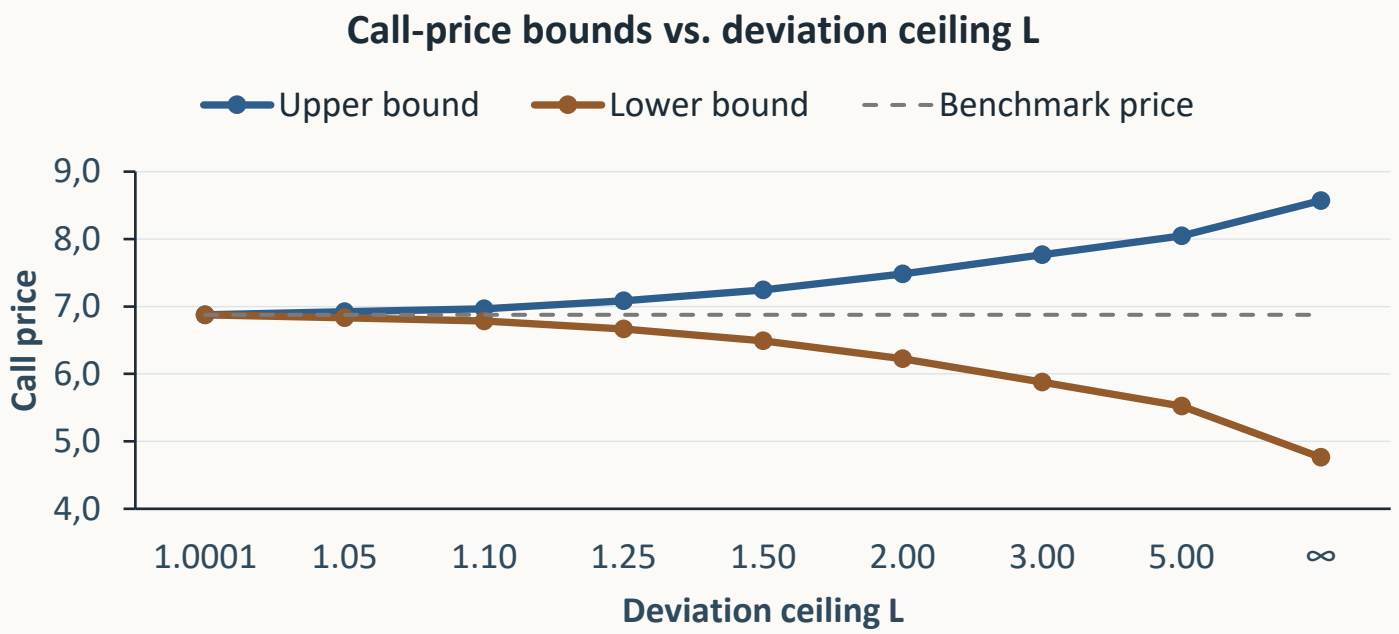

TAKEAWAY

The band interpolates smoothly — bounds widen monotonically from the benchmark price **6.88** toward the full no-arbitrage interval **[4.76, 8.57]** as $L \to \infty$, more valuation rules become admissible causing the price interval to widen until it reaches the full no arbitrage bounds:

Width grows from **0.0002** to **3.81**.

Intuition in the call example

- Think of the benchmark price 6.88 as our starting estimate. The band determines how far we are willing to move away from the benchmark valuation rule.
- The call payoff is $z = (0, 0, 15),$ hence only state 3 matters.
- Therefore, $\langle m, z \rangle = 15\, p_3\, m_3,$ the option value depends entirely on how much weight the kernel puts on state 3.
- For each value of L, we solve two linear programs. The upper bound selects the admissible pricing kernel within the band that assigns the highest value to the call payoff, while the lower bound selects the admissible kernel assigning the lowest value. Put differently:
- The upper bound asks 'what is the most optimistic call price consistent with that tolerance?', while the lower bound asks 'what is the most pessimistic price?'.
- Increasing Lexpands the set of admissible valuation rules, so the optimistic and pessimistic assessments move farther apart.

Conclusion

What we showed

- Finite-state bridge via the duality $\langle m, z \rangle$ between payoffs and kernels
- Multiplicative band $r\, b_s \le m_s \le L\, r\, b_s$ keeps SDFs near the ray $[\, b \,]$
- Exact recovery: $[\, b \,]=[\, m^* \,] \Rightarrow M(L\,;\, b)=M_{BL}(L\,;\, m^*)$
- Span collapse consistent with Chambers–Färe exact pricing; reference governed by Hilbert distance
- The paper turns "which pricing kernels should we trust?" into a geometric question about distance from a benchmark ray.

Future work

- Extend the construction to **dynamic settings**
- Data-driven rules for **selecting benchmark rays**
- Compare **multiplicative vs. quadratic** kernel neighborhoods empirically

Thank you — comments and questions welcome