

Executable Robust Arbitrage Certificates on Prediction Sets: Attainment and Faithful Discretization*

Shaosai Huang[†]

September 5, 2026

Abstract

On a finite-calendar prediction set with a finite executable traded menu and bid–ask quoted claims, we distinguish inconsistency with the martingale class from executable arbitrage: a finite static portfolio and predictable strategy whose profit is positive on every declared path. On finite trees, a Farkas alternative gives exactly one of a calibrated martingale law or such a certificate, with an attained margin, and an exact carry identity converts fixing-date marks into payment ledgers. On compact continuum-state prediction sets, under compactness, continuity, rectangularity, and saturation, the key quantitative condition is a uniform relative-interior ball of radius κ_n in each one-step increment hull. It yields attained one-step duality with an optimal slope bounded by $2\|f\|_\infty/\kappa_n$, hence attained pointwise super- and subhedging duality, bounded Borel optimizers, optional decomposition, a closed gain cone with the martingale class as exact polar, and finite certificates for uncountable quote books. Across families of pointwise-interior one-step geometries, a common balance bound makes the uniform radius equivalent to uniformly bounded least-norm optimizers; without it the converse fails. Martingale cubature can recover the martingale class and every robust value while tree radii collapse and optimal slopes diverge. Adding at most $2d_n(d_n + 1)$ successors per node restores radius $\kappa_n/\sqrt{d_n}$, value and margin errors admit quantitative bounds, and prefix-compatible shadow maps lift finite-tree inequalities with an explicit buffer, yielding continuum certificates when it is positive.

Keywords. Robust superhedging; prediction sets; bid–ask spreads; arbitrage certificates; optional decomposition; martingale cubature; discretization of martingale classes.

JEL classification. G13, G12, C61.

MSC 2020. 91G20, 60G42, 90C05, 49N15, 65D32.

1 Introduction

When market quotes cannot be matched by any model in a chosen library, the quotes have rejected the library; they have not yet produced an arbitrage. This distinction is the starting point of robust pricing on prediction sets, a formulation that descends from statistical prediction regions (Mykland, 2000, 2003), through the beliefs of Hou and Obłój (2018), to pathwise duality on prediction sets (Bartl et al., 2020), and, alongside it, the model-free fundamental theorem of Acciaio et al. (2016) for markets with statically traded options. One declares the set of price paths one takes seriously and asks what the quotes imply for every martingale law on it. Turning a rejection into an arbitrage

*Working paper. Comments welcome.

[†]Kspectra Research Inc., Toronto, ON, M2N 0G3, Canada. Email: arthur.foxie.huang@kspectra.ai.

then needs two further steps: the inconsistency must hold against every martingale law on the declared set, not only against the preferred models, and it must be converted into an explicit trade, a finite static portfolio and a predictable strategy whose realized profit is positive on every declared path.

This paper sets out to supply both steps in the following setting. Fix finitely many trading dates, a prediction set Ω , a finite vector X of executable common-numeraire prices, and statically quoted claims with bid and ask sides; the same price vector is written S in the Bellman formulation below. The object sought is

$$F_{a^+, a^-}(\omega) + (H \cdot X)_N(\omega) - \pi(a^+, a^-) \geq \gamma > 0 \quad \text{for every } \omega \in \Omega, \quad (1)$$

where $a^\pm$ are finite static positions bought at ask and sold at bid, π is their actual acquisition cost, and H is a predictable strategy in the traded menu. Inequality (1) is a booked trade, not a separating functional or a positive expected-value gap.

Two settings are treated. On a finite prediction tree the problem is a linear programme, and the certificate comes straight out of a Farkas alternative. On a compact continuum of states, duality still gives the right value, but a value is not a hedge: nothing yet guarantees that the optimal hedge exists as a bounded strategy. What guarantees it is a uniform relative-interior radius: for every date n there is $\kappa_n > 0$ such that

$$\{z \in D_n : |z| \leq \kappa_n\} \subseteq \text{conv}\{\Delta S_n(h, e) : e \in \Gamma_n(h)\} \quad \text{for every history } h, \quad (2)$$

where D_n is the active increment space. Under this condition an optimal one-step hedge can be chosen with slope at most

$$|H| \leq \frac{2\|f\|_\infty}{\kappa_n}. \quad (3)$$

This bound is the engine of both the attainment theorem and the discretization analysis. The paper's central computational finding is that martingale cubature can reproduce every robust value while failing to preserve (2): prices converge, yet the optimal tree hedges blow up. A finite nodewise augmentation restores a uniform radius, and with an off-grid shadow estimate a positive buffered margin then yields an executable continuum certificate.

The main theorem, heuristically. Take a compact prediction set with continuously varying successor sets, a finite traded menu, and, as dual class, the full set of martingale laws on the declared set. Under the radius condition (2) — at every history the market can move at least κ_n in every active direction — robust pricing on the declared set is exact and executable. Superhedging and subhedging prices are attained by bounded hedges, not only by laws; a quote book inconsistent with every martingale law, finite or not, yields a finite trade with negative cost and positive payoff on every declared path; and computation is faithful once the cubature trees are augmented to keep the radius — values and margins converge, tree hedges can be kept bounded, and, under an off-grid shadow estimate, a tree certificate whose margin clears an explicit buffer is a continuum certificate. One bound carries all of it: the radius limits every one-step hedge to (3). The precise statements are [Corollary 4.3](#) and [Theorem 7.10](#).

1.1 The radius, attainment, and its precise scope

The mechanism is a one-step fact ([Lemma 6.1](#)). On a compact successor set, the ball (2) makes the one-step hedging problem dual to a martingale-kernel problem with both sides attained, and it caps the optimal slope by (3). Everything else follows by iterating this through the Bellman

recursion: the value functions are continuous, the optimizers are bounded and Borel, super- and subhedging duality holds pointwise with all extrema attained, the price process has a pathwise optional decomposition, and continuous ε -hedges exist (Theorem 6.6). Because the ball lets every successor carry positive mass under some martingale kernel, every declared path is charged by some law in $\mathcal{M}_\Omega$ (Lemma 5.3), so pointwise statements need no efficient-set qualification; the gain cone of continuous payoffs is closed, with $\mathcal{M}_\Omega$ as its exact polar (Corollary 6.7); and compactness reduces an infeasible quote book of any size to a finite infeasible sub-book, hence to a finite certificate (Theorem 6.8). The decomposition is the compact-state descendant of the classical one (Kramkov, 1996; Föllmer and Kramkov, 1997); for the dominated discrete-time theory see Föllmer and Schied (2025).

How much is the radius really buying? Two things, best kept apart. It is sufficient and cannot be dropped: with zero on the boundary of the increment hull, the one-period payoff $f(y) = -\sqrt{y}$ has a finite value but no finite attaining slope (Section 8.2). It is necessary only up to balance: for families of one-step geometries in which opposite moves come in comparable size, uniformly bounded least-norm hedges force a common positive radius (Theorem 6.4), while without such balance an asymmetric two-point family has collapsing radii and bounded hedges, and the two-sided quantity of Proposition 6.2 is what actually controls them. Inside a single compact market with continuous successor sets none of this subtlety arises, because pointwise interiority already makes the radius continuous and hence uniformly positive (Remark 6.5); uniformity is a genuine extra requirement only across a family of markets, which is exactly the situation of a discretization sequence.

This sharpens the existence theory rather than replacing it. Robust duality with measurable strategies holds under far weaker analytic hypotheses (Bouchard and Nutz, 2015; Nutz, 2014; Burzoni et al., 2017, 2019; Oblój and Wiesel, 2021; Cheridito et al., 2017); dynamic characterizations of the robust price are given by Carassus et al. (2019); that dual attainment can fail, or needs relaxation, is shown in (Beiglböck et al., 2013, 2017); full support and efficient sets are studied by Burzoni et al. (2016); and Riedel (2015) attains both sides in one period on a compact space. What the radius adds is a dynamic completion that is bounded, uniform over histories, and valid on every path, and that can be tracked through finite-tree computation.

1.2 Faithful discretization

In practice a robust certificate is computed on a finite tree, so the question is what survives the passage to the tree. Raw-history martingale cubature (Lemma 7.1), a Tchakaloff-type construction (Tchakaloff, 1957; Bayer and Teichmann, 2006), produces nested finite subtrees $\Omega_m \subseteq \Omega$ whose full martingale classes $\mathcal{M}_m$ converge to $\mathcal{M}_\Omega$ in the weak Hausdorff metric (Proposition 7.2); every robust value therefore converges, in line with the value-level discretization literature (Guo and Oblój, 2019; Dolinsky and Soner, 2013, 2014a).

Values, however, are not hedges. Proposition 7.3 gives a two-step market with continuum radius one and cubature trees that recover the whole martingale class while some of their local radii shrink like 2^{-2j} ; for one fixed continuous payoff the tree-optimal slopes then diverge. Law and value convergence thus guarantee neither bounded tree hedges nor validity off the grid: the lattice prices correctly and hedges wrongly.

The repair is local and geometric. At each node of active dimension $d_n \geq 1$, adjoining at most $2d_n(d_n + 1)$ allowed successors puts a cross-polytope, hence a ball of radius $\kappa_n/\sqrt{d_n}$, inside the increment hull (Lemma 7.4 and Proposition 7.5). On the augmented trees, values converge for every continuous payoff; for each fixed noncrossed finite quote list the separation margins converge and bounded Bellman-optimal holdings can be selected (Theorem 7.6 and Corollary 7.9); the value

and margin errors are $O(1/m)$ on a finite span of the metrizing test functions, with approximation bounds for general continuous payoffs (Proposition 7.7); and under the prefix-compatible shadow estimates of Proposition 7.8 a finite-tree certificate lifts to every continuum path with the explicit buffer $\gamma_m - \sqrt{q} \xi_m - 2NL\eta_m$, so it is a continuum certificate whenever that buffer is positive. Convergence of the hedge ratios themselves is not claimed; what is claimed is that the certificate — its margin, its cost, and its validity on every declared path — survives.

1.3 The executable certificate

On a finite tree no radius condition is needed: with finitely many successors the local envelope is polyhedral, so whenever the barycentric constraints are feasible a finite supporting slope exists. The Farkas alternative (Theorem 3.4) then says that exactly one of two things exists, a calibrated martingale law in the bid–ask polytope or an executable certificate. Its nonnegative dual variables are the long-at-ask and short-at-bid quantities, and its free node multipliers are a predictable strategy. The certificate can be normalized, the compact margin programme (25) attains a numerical inconsistency margin, and when the martingale polytope is nonempty finite-tree subhedging is attained on both sides (Proposition 3.5).

Executability also requires trading the contractual payoff rather than an equal-expectation mark. The finite-leg carry identity (Lemma 4.1) and its invariance (Proposition 4.2) show when fixing-date marks and payment ledgers have identical feasibility and static cost, with an explicit adjustment of the predictable strategy; likewise, a separating potential must be reduced to finitely many actually quoted legs, whether masses, forwards, calls, or puts. These are the ingredients that turn a separating value into the trade (1).

The transaction-cost and robust-friction literature establishes consistent price systems and fundamental alternatives in quasi-sure, pointwise, and model-free forms (Bouchard and Nutz, 2016; Bayraktar and Zhang, 2016; Bouchard et al., 2019; Bayraktar and Burzoni, 2020; Burzoni, 2016; Dolinsky and Soner, 2014b), building on the classical theory (Jouini and Kallal, 1995; Schachermayer, 2004; Kabanov and Safarian, 2009; Pennanen, 2011); finite-tree pricing and hedging algorithms are studied by Roux and Zastawniak (2009); static quote consistency and numerical repair are treated in (Davis and Hobson, 2007; Cousot, 2007; Gerhold and Gülüm, 2020; Cohen et al., 2020). What is new here is the finite, sign-structured, pathwise certificate with its attained margin and contractual carry, and its survival under discretization.

1.4 Contributions, scope, and organization

The tools used below are standard: finite-dimensional linear duality, the maximum theorem, measurable selection, Carathéodory and Tchakaloff cubature, and Hausdorff convergence. The contributions are three.

1. The uniform interior radius is identified as the mechanism behind bounded, attained, pointwise hedging, and its scope is made precise: across families of one-step geometries a common balance constant makes it equivalent to uniformly bounded hedges, and without balance a two-sided slope quantity takes over (Section 6).
2. Martingale cubature is shown to recover the full martingale class while letting local radii collapse; a finite nodewise augmentation restores the radius and yields uniform hedge bounds together with quantitative value, margin, and conditional certificate-lifting guarantees (Section 7).
3. Finite-tree Farkas duality, compact-state robust subhedging, finite-sub-book extraction, and exact fixing-to-payment carry are assembled into a pipeline from quote infeasibility to an

executable certificate (Sections 3 to 6).

The regime is compact and finite-calendar, with a finite dynamically traded menu; unbounded prices need a separate localization or truncation argument. Rectangularity is assumed rather than inferred, since nonrectangular restrictions generate untraded dual multipliers, the market analogue of the robust-control distinction between a law class and its rectangular hull (Epstein and Schneider, 2003; Iyengar, 2005; Nilim and El Ghaoui, 2005). Continuous-time and growing-menu extensions need additional closedness and stochastic-integral stability; single-reference-law stability results include (Kardaras, 2024; De Donno and Pratelli, 2005; Alberti and Karbach, 2026), and the nondominated growing-menu problem is open. The results are stated abstractly, with interest-rate contracts entering only as worked instances of the carry identity, and the examples in Section 8 delimit the present claims.

Section 2 fixes the market and certificate conventions. Sections 3 and 4 give the finite-tree alternative and carry identity. Sections 5 and 6 establish the compact-state martingale class, attained duality, optional decomposition, and finite executable certificates. Section 7 proves cubature recovery, radius collapse, augmentation, and lifting, and includes the finite-level separator-to-trade principle. Section 8 records the sharp boundaries and open questions, and Section 9 closes the argument. Contract-level carry examples are collected in appendix A.

2 Market, trading, and certificate conventions

This section fixes, once and exactly, the objects every later statement is made of: the finite master calendar and the canonical gains of a finite executable menu; the statically quoted claims with their bid–ask accounting and the definition of an executable robust arbitrage certificate; and the carry interface through which fixing-date expressions are attached to payment-date ledgers. Nothing here is deep; the point is that every later claim of executability is checked against these conventions and no others.

2.1 Finite master calendar and canonical gains

Fix deterministic dates

$$0 = t_0 < t_1 < \dots < t_N = T.$$

Every fixing publication, exercise, contractual carry, delivery, and payment date used by the finite certificate belongs to this calendar. The state at t_i is a history $h_i = (x_0, \dots, x_i, a_0, \dots, a_i)$, where a is the revealed shared fixing/settlement tape and $x_i \in \mathbb{R}^d$ is a finite vector of common-terminal-numeraire primitive prices. It contains every state-price bond, gross-coupon price, state-price annuity, or linear switch claim that the certificate trades. Redundant nonlinear rates are derived from these primitive coordinates and are not traded as separate assets.

At a publication time, the history used for a trade is the marked right-hand (post-publication) state. Equivalently, a market with both a pre-publication and a post-publication decision splits them into two ordered nodes. Thus every fixing-dependent coefficient used in a holding selected at t_i is already $\mathcal{F}_{t_i}$ -measurable; the strategy never sees an unpublished fixing.

For a predictable strategy $H = (H_i)_{i=0}^{N-1}$, with $H_i(h_i) \in \mathbb{R}^d$, define the discounted gain

$$(H \cdot X)_N := \sum_{i=0}^{N-1} H_i(h_i) \cdot (x_{i+1} - x_i). \quad (4)$$

The reference asset corresponding to one unit of the selected terminal numeraire is the constant one after normalization. It need not be the money-market account. Thus a cash coordinate z and

the gain (4) are already expressed in common terminal-numeraire units. On a finite tree every predictable strategy is admissible. On a compact continuum-state prediction set we shall use the bounded Borel class, and its continuous subclass, declared in [Section 5](#); no unstated doubling strategy is allowed.

2.2 Static claims and bid–ask accounting

Let $f^1, \dots, f^q$ be the terminal-numeraire payoffs of the finitely many vanilla or linear-switch claims actually used in a certificate. Their market bid and ask prices are finite:

$$-\infty < \ell_k, u_k < \infty.$$

No ordering is imposed in the finite-tree fundamental alternative: if $\ell_k > u_k$, the quote is crossed and the alternative returns the direct buy-at-ask/sell-at-bid round trip. Results that use a nonempty quote box, including the compact-state separation theorem, explicitly impose $\ell_k \leq u_k$. These are normalized prices. If the currency quotes are $\ell_k^{\text{ccy}}, u_k^{\text{ccy}}$ and the common initial state-price scale is ζ_0 , then $\ell_k = \zeta_0 \ell_k^{\text{ccy}}$ and $u_k = \zeta_0 u_k^{\text{ccy}}$. The long and short quantities are separate vectors $a^+, a^- \in \mathbb{R}_+^q$. The initial acquisition cost and terminal payoff are

$$\pi(a^+, a^-) := \sum_{k=1}^q (a_k^+ u_k - a_k^- \ell_k), \quad (5)$$

$$F_{a^+, a^-} := \sum_{k=1}^q (a_k^+ - a_k^-) f^k. \quad (6)$$

This convention is load-bearing: buying costs the ask and shorting receives the bid. Replacing a^+, a^- by one signed coefficient and pricing it at a single mid would prove a different and economically invalid theorem.

Definition 2.1 (Calibrated pricing law). A law Q on the declared prediction set is calibrated to the finite market if X is a true Q -martingale in the raw canonical filtration and

$$\ell_k \leq E_Q f^k \leq u_k, \quad k = 1, \dots, q. \quad (7)$$

The set of all such laws is denoted by $\mathcal{M}(\ell, u)$.

If $Q \in \mathcal{M}(\ell, u)$, then for every bounded predictable H ,

$$E_Q[F_{a^+, a^-} + (H \cdot X)_N] \leq \pi(a^+, a^-). \quad (8)$$

Indeed $E_Q(H \cdot X)_N = 0$, while each long expectation is at most its ask and each short expectation is at least its bid. Consequently, the following pathwise object rules out every calibrated law.

Definition 2.2 (Executable robust arbitrage certificate). An executable certificate is a triple (a^+, a^-, H) and a number $\gamma > 0$ such that

$$F_{a^+, a^-}(\omega) + (H \cdot X)_N(\omega) - \pi(a^+, a^-) \geq \gamma \quad \text{for every admissible path } \omega. \quad (9)$$

It is a strict, pointwise terminal-numeraire profit after its actual initial bid–ask cost. Equivalently, after adding a position in the constant account, one may write the same opportunity as a strictly negative initial cost (a positive receipt) and a nonnegative terminal payoff; the exact cash conversion is displayed in (20)–(21). We keep (9) throughout to avoid sign ambiguity.

2.3 The carry interface

Expectation-level separators of the kind discussed in [Section 7.6](#) produce finite piecewise-linear fixing-date expressions of the form

$$g = \sum_{r=1}^m \alpha_r W_{a_r} \phi_r(Y_{a_r}), \quad (10)$$

where every ϕ_r is a finite linear combination of constants, the identity, and call payoffs $(y - K)^+$. The corresponding market cost is the same linear combination of observed caplet/swaption prices, using ask for positive holdings and bid for negative holdings.

Definition 2.3 (Exact carry-admissibility). A fixing expression g is exactly carry-admissible if there are an actually traded terminal-numeraire payoff $\tilde{g}$ and a predictable zero-initial-cost strategy H^{car} in the declared traded coordinates such that

$$\tilde{g} - g = (H^{\text{car}} \cdot X)_N \quad \text{pathwise on the prediction set,} \quad (11)$$

and the static acquisition cost of $\tilde{g}$ is the quoted cost of the finite vanilla portfolio representing g . All denominators and settlement weights must stay in the contract domain along every admissible path.

The caplet, physical swaption, declared linear cross-CSA switch, and declared cash-settlement instances are proved in [Appendix A](#). The abstract definition is retained because no pathwise trading theorem can manufacture a missing settlement asset or turn an untraded continuum strike potential into a finite portfolio.

Remark 2.4 (Raw filtration). All martingale and predictability statements use the raw canonical filtration. Law-dependent completions are taken only after a law is fixed. This is required both for convex mixing of pricing laws and for a single pathwise meaning of the strategy H .

3 Exact finite-tree trading certificates

This section turns finite quote infeasibility into an executable trading statement. There is no limiting argument here: the prediction tree, the traded gain coordinates, and the quote panel are all finite. The result is therefore a theorem of linear alternatives. The economically important point is that its dual variables have exactly the signs of a bid–ask portfolio and exactly the measurability of a predictable strategy.

3.1 The canonical finite prediction tree

Fix a master calendar

$$0 = t_0 < t_1 < \dots < t_N = T_\star$$

containing every determination, exercise, carry, and payment date in the finite panel under consideration. Let $\mathbb{T} = \bigsqcup_{i=0}^N \mathbb{T}_i$ be a finite rooted tree. Its root is $o \in \mathbb{T}_0$, every $v \in \mathbb{T}_i$, $i < N$, has a nonempty finite child set $\text{ch}(v) \subset \mathbb{T}_{i+1}$, and every nonroot node has one predecessor. The leaves

$$\Omega_{\mathbb{T}} := \mathbb{T}_N$$

are the declared prediction paths. If $\omega \in \Omega_{\mathbb{T}}$, write $v_i(\omega)$ for its unique level- i ancestor.

Let

$$S_i(v) = (S_i^1(v), \dots, S_i^d(v)) \in \mathbb{R}^d, \quad v \in \mathbb{T}_i,$$

be a finite panel of *executable carried gain coordinates*, expressed in one common reference state-price unit. A matured coordinate is stopped, so keeping it in the vector only adds zero future increments. The constant unit account is denoted by $S^0 \equiv 1$ and is handled as the cash coordinate below. Merely being a canonical state variable is not enough for inclusion in S : each component of S must be the value process of a traded linear claim, including its contractual collateral gain stream.

A predictable simple strategy is a family

$$H = (H_i(v))_{0 \leq i < N, v \in \mathbb{T}_i}, \quad H_i(v) \in \mathbb{R}^d.$$

Its terminal gain on a leaf is

$$(H \cdot S)_N(\omega) = \sum_{i=0}^{N-1} H_i(v_i(\omega)) \cdot (S_{i+1}(v_{i+1}(\omega)) - S_i(v_i(\omega))). \quad (12)$$

All such strategies are bounded because the node set is finite. In particular, there is no hidden admissibility or stochastic-integrability issue in the finite-tree setting.

For $v \in \mathbb{T}_i$, $i < N$, and $1 \leq k \leq d$, define the row payoff

$$d_{v,k}(\omega) = \mathbf{1}_{\{v_i(\omega)=v\}} (S_{i+1}^k(v_{i+1}(\omega)) - S_i^k(v)). \quad (13)$$

After choosing an ordering of nodes, assets, and leaves, let D be the matrix with entries $D_{(v,k),\omega} = d_{v,k}(\omega)$. If h is the vector formed by the node holdings $H_i^k(v)$, then

$$((H \cdot S)_N(\omega))_\omega = D^\top h. \quad (14)$$

A law is a vector $q = (q_\omega)_{\omega \in \Omega_{\mathbb{T}}}$. Put

$$\mathcal{Q}_{\mathbb{T}} = \{q \in \mathbb{R}_+^{\Omega_{\mathbb{T}}} : \mathbf{1}^\top q = 1, Dq = 0\}. \quad (15)$$

For a node v , let $q(v)$ be the mass of its descendant leaves. If $q(v) > 0$, the row equations belonging to v say precisely

$$\sum_{w \in \text{ch}(v)} \frac{q(w)}{q(v)} S_{i+1}(w) = S_i(v).$$

If $q(v) = 0$, they reduce to $0 = 0$. Thus (15) is equivalent to martingality of the traded panel in its raw tree filtration, without imposing an artificial conditional law at a null node.

Proposition 3.1 (Local barycentric alternative). *At a nonterminal node $v \in \mathbb{T}_i$, $i < N$, exactly one of the following holds:*

1.

$$S_i(v) \in \text{conv}\{S_{i+1}(w) : w \in \text{ch}(v)\};$$

2. *there are $a_v \in \mathbb{R}^d$ and $\delta_v > 0$ such that*

$$a_v \cdot (S_{i+1}(w) - S_i(v)) \geq \delta_v \quad \text{for every } w \in \text{ch}(v). \quad (16)$$

In the second case, holding a_v for the one step after v is a zero-cost conditional dynamic arbitrage on the subtree rooted at v . If $v = o$, then taking cash $c = -\delta_v/2$ makes it a global negative-cost certificate with terminal payoff at least $\delta_v/2$.

Proof. The first set is a nonempty compact convex polytope. If it does not contain $S_i(v)$, strict separation gives a_v whose value on every successor exceeds its value at $S_i(v)$; finiteness of the child set makes the minimum gap δ_v strictly positive. Conversely, (16) is incompatible with any convex combination having barycenter $S_i(v)$. $\square$

A failed barycentric condition at a nonroot node is only a conditional arbitrage unless reaching that node is itself forced or hedgeable. A global martingale law may assign that node zero mass. The compact-state theory of Section 5 therefore either prunes such nodes or assumes local barycentric feasibility at every history in its rectangular prediction set.

Let $g^1, \dots, g^J : \Omega_{\mathbb{T}} \rightarrow \mathbb{R}$ be the actual common-unit terminal ledger payoffs of the quoted claims. Write

$$G_{j,\omega} = g^j(\omega), \quad \underline{p}, \bar{p} \in \mathbb{R}^J$$

for their bid and ask vectors, respectively. The Farkas theorem does not require them to be noncrossed: if some $\underline{p}_j > \bar{p}_j$, it returns the immediate buy-at-ask, sell-at-bid round trip. The calibrated-law polytope is

$$\mathcal{Q}_{\mathbb{T}}^{\text{ba}} = \{q \in \mathcal{Q}_{\mathbb{T}} : \underline{p} \leq Gq \leq \bar{p}\}. \quad (17)$$

All inequalities between vectors are componentwise.

For long and short quantities $\vartheta^+, \vartheta^- \in \mathbb{R}_+^J$, the executable static payoff and its acquisition cost are

$$F_{\vartheta} = \sum_{j=1}^J (\vartheta_j^+ - \vartheta_j^-) g^j, \quad (18)$$

$$\pi^{\text{ba}}(\vartheta^+, \vartheta^-) = \bar{p} \cdot \vartheta^+ - \underline{p} \cdot \vartheta^-. \quad (19)$$

A positive coefficient is bought at the ask; a negative coefficient is shorted for the bid. This sign convention will be recovered, rather than inserted after the duality calculation.

Definition 3.2 (Finite-tree executable certificate). An executable certificate consists of $c \in \mathbb{R}$, a predictable strategy H , and static quantities $\vartheta^+, \vartheta^- \in \mathbb{R}_+^J$, such that

$$\mathbf{C}_0 := c + \bar{p} \cdot \vartheta^+ - \underline{p} \cdot \vartheta^- < 0, \quad (20)$$

$$\mathbf{V}_N(\omega) := c + (H \cdot S)_N(\omega) + F_{\vartheta}(\omega) \geq 0 \quad \text{for every } \omega \in \Omega_{\mathbb{T}}. \quad (21)$$

The inequality is pathwise on the declared tree, not merely almost surely under one selected model.

This is equivalent to the profit-after-cost convention in (9). Indeed, if $\mathbf{C}_0 = c + \pi^{\text{ba}} < 0$, then

$$F_{\vartheta} + (H \cdot S)_N - \pi^{\text{ba}} = \mathbf{V}_N - \mathbf{C}_0 \geq -\mathbf{C}_0 > 0.$$

Conversely, if the left side is at least $\gamma > 0$, choosing $c = -\pi^{\text{ba}} - \gamma/2$ gives initial cost $-\gamma/2$ and terminal payoff at least $\gamma/2$.

The certificate already is an arbitrage under the negative-cost convention. Under a zero-initial-capital convention, buy $-\mathbf{C}_0 > 0$ units of the constant account with the initial receipt. The resulting terminal payoff is $\mathbf{V}_N - \mathbf{C}_0 > 0$ on every prediction path.

3.2 The exact LP/Farkas alternative

With the tree, the executable panel, and the certificate class fixed, the fundamental alternative is a statement about one finite linear system. We first record the mixed Farkas lemma in exactly the form needed — equality rows for the martingale constraints, inequality rows for the bid–ask sides — and then read its dual variables as a trade.

Lemma 3.3 (Mixed Farkas alternative). *Let $A \in \mathbb{R}^{m \times L}$, $B \in \mathbb{R}^{n \times L}$, $b \in \mathbb{R}^m$, and $d \in \mathbb{R}^n$. Exactly one of the following systems is feasible:*

$$\begin{aligned} x \geq 0, \quad Ax = b, \quad Bx \leq d; \\ y \in \mathbb{R}^m, \quad z \in \mathbb{R}_+^n, \quad A^\top y + B^\top z \geq 0, \quad b \cdot y + d \cdot z < 0. \end{aligned}$$

Proof. Introduce a slack $s \geq 0$. The first system is feasible exactly when (b, d) belongs to the finitely generated, hence closed, convex cone

$$\mathcal{C} = \{(Ax, Bx + s) : x \geq 0, s \geq 0\}.$$

If it does not, strong separation of the point from $\mathcal{C}$, followed by positive rescaling, gives (y, z) such that

$$b \cdot y + d \cdot z < 0 \leq (Ax) \cdot y + (Bx + s) \cdot z \quad (x, s \geq 0).$$

Varying s gives $z \geq 0$, and varying x gives $A^\top y + B^\top z \geq 0$. Conversely, multiplying a feasible first system by such y, z yields a number which is simultaneously nonnegative and strictly negative. This is the standard mixed Farkas lemma; compare [Rockafellar \(1970, Section 22\)](#). $\square$

Theorem 3.4 (Finite-tree bid–ask fundamental alternative). *Exactly one of the following systems is feasible:*

1. *there is a calibrated martingale law $q \in \mathcal{Q}_{\mathbb{T}}^{\text{ba}}$;*
2. *there is an executable certificate in the sense of [Definition 3.2](#).*

If the second system is feasible, a certificate can be normalized so that $C_0 = -1$. Hence the strict alternative and its certificate are attained by a finite linear program.

Proof. The calibrated-law problem is the mixed linear system

$$\begin{aligned} q &\geq 0, \\ \begin{bmatrix} \mathbf{1}^\top \\ D \end{bmatrix} q &= \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \\ \begin{bmatrix} G \\ -G \end{bmatrix} q &\leq \begin{bmatrix} \bar{p} \\ -\underline{p} \end{bmatrix}. \end{aligned} \tag{22}$$

The mixed equality–inequality alternative [Lemma 3.3](#) says that (22) is infeasible if and only if there are free multipliers $c \in \mathbb{R}$, $h \in \mathbb{R}^{\text{rows}(D)}$ and nonnegative multipliers $\vartheta^+, \vartheta^- \in \mathbb{R}_+^J$ satisfying

$$c\mathbf{1} + D^\top h + G^\top (\vartheta^+ - \vartheta^-) \geq 0, \tag{23}$$

$$c + \bar{p} \cdot \vartheta^+ - \underline{p} \cdot \vartheta^- < 0. \tag{24}$$

For completeness, the easy direction follows by multiplying (23) by a feasible q : its left side has nonnegative q -average, whereas the equality and inequality constraints in (22) make that average

no larger than the strictly negative quantity in (24). The converse is precisely the separating-hyperplane proof of Farkas' lemma applied to the polyhedral cone generated by the columns of the augmented constraint matrix.

Use (14) to read h as a predictable strategy. Then (23) is (21), and (24) is (20). Thus every Farkas multiplier is economically executable with the correct side of each quote. Conversely, the coefficients of every certificate satisfy the two Farkas inequalities. Scaling by the absolute value of the strictly negative cost gives $C_0 = -1$.

The primal feasible region is a closed subset of the probability simplex and hence a compact polytope. On the certificate side, existence is already given by Farkas' lemma. If one wants a compact quantitative LP, split the free variables into positive and negative parts and solve

$$\begin{aligned} \max \quad & \gamma \\ \text{subject to} \quad & c\mathbf{1} + D^\top h + G^\top(\vartheta^+ - \vartheta^-) \geq 0, \\ & c + \bar{p} \cdot \vartheta^+ - \underline{p} \cdot \vartheta^- + \gamma \leq 0, \\ & |c| + \|h\|_1 + \mathbf{1} \cdot (\vartheta^+ + \vartheta^-) \leq 1, \\ & \gamma \geq 0, \quad \vartheta^\pm \geq 0. \end{aligned} \tag{25}$$

After the standard linearization of absolute values, its feasible set is compact. Its optimum is attained, is zero when $\mathcal{Q}_{\mathbb{T}}^{\text{ba}} \neq \emptyset$, and is strictly positive when that polytope is empty: normalize any nonzero Farkas certificate in the displayed ℓ^1 -gauge. This supplies an attained numerical inconsistency margin as well as a feasibility certificate. $\square$

If the quotes are not crossed, simultaneous long and short positions in one claim are unnecessary. Indeed, subtracting $m_j = \min(\vartheta_j^+, \vartheta_j^-)$ from both sides leaves the payoff unchanged and changes the cost by $-m_j(\bar{p}_j - \underline{p}_j) \leq 0$. A certificate can therefore be chosen with $\vartheta_j^+ \vartheta_j^- = 0$ for every j .

3.3 Finite robust subhedging and the constructive separator route

The preceding theorem also decomposes into the two operations needed to turn an expectation-level separation into a trade: static separation followed by dynamic subhedging.

Proposition 3.5 (Finite-tree robust subhedging). *Suppose $\mathcal{Q}_{\mathbb{T}} \neq \emptyset$. For every terminal payoff $F : \Omega_{\mathbb{T}} \rightarrow \mathbb{R}$,*

$$\min_{q \in \mathcal{Q}_{\mathbb{T}}} E_q[F] = \max \{z \in \mathbb{R} : \exists H \text{ predictable, } z + (H \cdot S)_N \leq F \text{ on } \Omega_{\mathbb{T}}\}. \tag{26}$$

Both extrema are attained.

Proof. In matrix form, the two sides are the primal–dual pair

$$\begin{aligned} \min_{q \geq 0} \quad & F^\top q & \text{subject to} \quad & \mathbf{1}^\top q = 1, \quad Dq = 0, \\ \max_{z, h} \quad & z & \text{subject to} \quad & z\mathbf{1} + D^\top h \leq F. \end{aligned}$$

The primal is feasible by assumption and has a compact feasible polytope. Its value is finite. Finite-dimensional linear-programming duality, the objective-value form of Lemma 3.3, gives equality and dual attainment. The vector h is a predictable strategy by (14). $\square$

At a fixed finite node, no relative-interior hypothesis is needed for this duality. Indeed, if the successor values are $s_1, \dots, s_m$, the local lower-envelope problem

$$\min \left\{ \sum_{\ell=1}^m \lambda_\ell f_\ell : \lambda \geq 0, \quad \sum_{\ell} \lambda_\ell = 1, \quad \sum_{\ell} \lambda_\ell s_\ell = x \right\}$$

has the dual

$$\max_{\alpha, h} \{ \alpha + h \cdot x : \alpha + h \cdot s_\ell \leq f_\ell \text{ for every } \ell \}.$$

Whenever the barycentric constraint is feasible, both finite LPs attain. Equivalently, the lower convex envelope is polyhedral and has at least one finite supporting slope at x , even when x lies on the relative boundary of the successor hull. The continuum difficulty is not existence of a slope at one node; it is obtaining a measurable, uniformly admissible selection as the history and successor set vary. That is why the compact-state theory needs additional continuity and slope-control hypotheses.

Assume now that $\mathcal{Q}_{\mathbb{T}} \neq \emptyset$, every quote is noncrossed, and put

$$\mathcal{K}_{\mathbb{T}} := \{Gq : q \in \mathcal{Q}_{\mathbb{T}}\}, \quad \mathcal{B} := [\underline{p}, \bar{p}].$$

Both are compact convex sets, and $\mathcal{Q}_{\mathbb{T}}^{\text{ba}} \neq \emptyset$ exactly when $\mathcal{K}_{\mathbb{T}} \cap \mathcal{B} \neq \emptyset$. If they are disjoint, strict finite-dimensional separation, with its sign reversed if necessary, gives a vector $\vartheta \in \mathbb{R}^J$ such that

$$\underbrace{\inf_{q \in \mathcal{Q}_{\mathbb{T}}} E_q \left[\sum_j \vartheta_j g^j \right]}_{\text{robust lower model value}} > \underbrace{\sup_{p \in \mathcal{B}} \vartheta \cdot p}_{\bar{p} \cdot \vartheta^+ - \underline{p} \cdot \vartheta^-} = \pi^{\text{ba}}(\vartheta^+, \vartheta^-). \quad (27)$$

The displayed support-function identity fixes the orientation:

$$\sup_{p \in \mathcal{B}} \vartheta \cdot p = \bar{p} \cdot \vartheta^+ - \underline{p} \cdot \vartheta^-.$$

Thus the separated claim is acquired long at asks and short at bids. Choose z strictly between the two sides. By [Proposition 3.5](#), an attained predictable H satisfies

$$z + (H \cdot S)_N \leq \sum_j \vartheta_j g^j \quad \text{on } \Omega_{\mathbb{T}}.$$

Buy the static portfolio and short this subhedge. Its initial cost is

$$\pi^{\text{ba}}(\vartheta^+, \vartheta^-) - z < 0,$$

while its terminal payoff is

$$\sum_j \vartheta_j g^j - z - (H \cdot S)_N \geq 0. \quad (28)$$

This is the conversion of a separating value into a trade described in the introduction. If $\mathcal{Q}_{\mathbb{T}} = \emptyset$, the direct Farkas theorem remains valid with $J = 0$ and gives a purely cash/dynamic global certificate; the robust lower value in (27) should not be used with an empty model set.

The theorem uses the full martingale class on the declared tree. A restriction to a preferred subclass introduces dual multipliers which need not be traded claims, and almost-sure nonnegativity under one law does not replace nonnegativity on every declared leaf. Likewise, a state coordinate cannot be placed in the gain vector merely because candidate laws make it a martingale: it must be executable. The explicit failures are collected in [Section 8](#).

4 Exact carry from fixing marks to payment ledgers

The finite potentials produced by an expectation-level separation are written naturally as fixing- or exercise-date state-price marks. A traded option, however, delivers its contractual payment ledger. Equality of their expectations under the model class would not by itself make one a hedge for the other. This section supplies the required pathwise self-financing identity.

4.1 Normalization and a finite-leg carry lemma

Fix a common initial state-price scale $\zeta_0 > 0$. For each declared linear collateral convention c and payment date $\tau = t_{n(\tau)}$ in the master calendar, let $Z_t^c(\tau)$ be the common-terminal-numeraire price at $t \leq \tau$ of one unit of payout currency at τ , and put $\zeta_t^c := Z_t^c(t)$. We assume $\zeta_0^c = \zeta_0$ and, whenever this payment claim is used in a carry, the stopped process

$$M_i^{c,\tau} := Z_{t_i \wedge \tau}^c(\tau), \quad M_{n(\tau)}^{c,\tau} = \zeta_\tau^c,$$

is a component of the executable panel S . These are pathwise state and tradability requirements; they do not select a reference law.

A currency bid or ask p^{ccy} is placed in the finite LP as

$$p^{\text{norm}} = \zeta_0 p^{\text{ccy}}. \quad (29)$$

Thus the model value of a common-unit payoff g is $E_Q g$, whereas its time-zero currency value is $\zeta_0^{-1} E_Q g$. For a payment C at date p under linear collateral convention c , the common-unit ledger payoff is

$$g = \zeta_p^c C. \quad (30)$$

This is one common state-price convention. For $c \neq 0$, it retains the relative collateral kernel and must not be replaced by the reference-collateral deflator. This normalization assumes the one-payout- currency setup. A foreign-currency payment is covered only after its executable FX conversion claim has been added to the traded state.

The following elementary telescoping identity is the load-bearing carry result.

Lemma 4.1 (Exact finite-leg carry). *For $1 \leq r \leq R$, let M^r be one component of the traded panel S , active from a determination level k through a payment level $n_r \geq k$ and stopped thereafter. Let χ_r be a finite $\mathcal{F}_{t_k}$ -measurable coefficient. Define the determination-date mark and the actual terminal ledger payoff by*

$$f = \sum_{r=1}^R \chi_r M_k^r, \quad g = \sum_{r=1}^R \chi_r M_{n_r}^r. \quad (31)$$

There is a bounded predictable carry strategy K , using only the displayed legs after determination, such that

$$g = f + (K \cdot S)_N \quad \text{on every leaf of the tree.} \quad (32)$$

Consequently $E_q g = E_q f$ for every $q \in \mathcal{Q}_{\mathbb{T}}$.

Proof. Level k is the marked right-hand, post-publication node at the determination event, so the coefficient χ_r is known before the holding on $[t_k, t_{k+1})$ is selected. If a venue or contract uses a different within-date order, the master calendar must contain separate pre-event and post-event

nodes. Hold χ_r units of M^r on every interval $i = k, \dots, n_r - 1$, and zero units outside those intervals. This is predictable. Pathwise telescoping gives

$$\chi_r M_{n_r}^r = \chi_r M_k^r + \sum_{i=k}^{n_r-1} \chi_r (M_{i+1}^r - M_i^r).$$

Sum over r to obtain (32). Taking expectation and using $Dq = 0$ gives the last assertion. Notice that the proof is pathwise and precedes any choice of law. $\square$

The lemma allows different payment dates and signed coefficients. A contract with optional collateral, nonlinear funding, or an untraded conversion factor is outside it unless that extra state is itself supplied with an executable linear carry decomposition.

Representative caplet, cash-settled swaption, physical-settlement, and cross-CSA instances are collected in [appendix A](#); they are applications of the pathwise identity, not additional hypotheses for the general theory.

4.2 Finite-strike executability

A finite piecewise-linear smile potential is executable through the same carry, but only once it has been reduced to actually quoted legs. Let $\phi : \mathbb{R} \rightarrow \mathbb{R}$ be continuous and piecewise affine, and let $K_1 < \dots < K_m$, $m \geq 0$, contain all slope changes. If $m \geq 1$, choose an anchor $K_0 < K_1$ in the first affine interval, let s_0 be the slope to the left of K_1 , and let s_r be the slope immediately to the right of K_r . If $m = 0$, choose any K_0 and let s_0 be the global affine slope. Then, with the sum interpreted as empty when $m = 0$,

$$\phi(y) = \phi(K_0) + s_0(y - K_0) + \sum_{r=1}^m (s_r - s_{r-1})(y - K_r)^+. \quad (33)$$

Therefore

$$W\phi(Y) = (\phi(K_0) - s_0 K_0)W + s_0 WY + \sum_{r=1}^m (s_r - s_{r-1})W(Y - K_r)^+. \quad (34)$$

The mass claim W , weighted-forward claim WY , and finitely many calls on the right must each be available with their own bid and ask and their declared carry. A quoted put may equivalently be retained as its own leg or reduced by put–call parity using these same mass and forward claims. For a variable-mass smile, omitting W would make the constant part of the potential nontradeable. If ϕ is constant in both tails, then $s_0 = 0$ and $\sum_r (s_r - s_{r-1}) = 0$, recovering the bounded-potential special case. The acquisition cost is the sum of the leg-level bid–ask costs; it is not a fictitious linear price assigned directly to ϕ .

Equation (34) is first an identity in the units of the potential’s affine primitive. Before execution, every term must be converted to the corresponding actual quote leg with the same deterministic notional, accrual, initial state-price or annuity mass, and ζ_0 scale. For example, if $W = \bar{A}_e / \bar{A}_0$, then the common-unit mark of a unit physical option is $\bar{A}_0 W (S_e - K)^+$; the traded quantity and its bid–ask cost must carry that factor. A fixed deterministic scale may equivalently be absorbed into the portfolio coefficient. A model-dependent mass may not: in that case the potential and the certificate use the unnormalized finite-measure weight and require the mass claim itself to be traded or finitely synthesized.

Proposition 4.2 (Carry invariance of the certificate cone). *Suppose every quoted claim has a fixing mark f^j , an actual ledger payoff g^j , and a predictable carry K^j satisfying*

$$g^j = f^j + (K^j \cdot S)_N \quad \text{pathwise.}$$

For any static net quantity $\vartheta_j = \vartheta_j^+ - \vartheta_j^-$ and any predictable H ,

$$\sum_j \vartheta_j g^j + (H \cdot S)_N = \sum_j \vartheta_j f^j + \left(\left(H + \sum_j \vartheta_j K^j \right) \cdot S \right)_N. \quad (35)$$

Here the last notation means the stochastic gain of the strategy in parentheses. Thus the actual-ledger and fixing-mark Farkas systems have the same feasibility status and the same static acquisition cost; their dynamic strategies differ by the explicit carry portfolio.

Proof. Multiply each pathwise carry identity by ϑ_j , sum over j , and use linearity of finite simple gains. Predictability is preserved because a finite linear combination of the K^j is predictable. $\square$

An equality of expectations without the pathwise identity is insufficient. On a two-leaf one-step space with no risky traded asset, let $f \equiv 0$ and $g(-) = -1, g(+) = 1$. They have equal expectation under the equal-weight law, but no self-financing strategy converts f into g . Replacing an actual ledger by a fixing mark solely on the strength of this expectation identity would therefore manufacture a false trading certificate.

Corollary 4.3 (Finite-tree certificate invariance under exact carry). *Suppose every quoted claim on the finite tree has a finite fixing mark, an actual payment ledger, and an exact predictable carry as in Lemma 4.1. Then calibrated-law feasibility and executable actual-ledger certificates obey the same strong alternative as in Theorem 3.4. Replacing marks by ledgers preserves bid–ask cost and the pathwise margin, and changes the dynamic strategy only by the explicit carry adjustment (35).*

Proof. Apply Theorem 3.4 to the fixing marks and then Proposition 4.2. The converse follows by reversing the same pathwise carry adjustment. $\square$

The hypotheses that must survive the compact-state lift are now visible. The continuum prediction set must be rectangular, or every nonrectangular restriction must have a traded dual; local martingale successor kernels must be nonempty; the value recursion and its argmin correspondence must be measurable; predictable slopes must obey an admissibility bound; actual contract ledgers must retain their pathwise finite-leg carries; and the terminal inequality must hold on the declared prediction set rather than only under one law. Compactness and continuity replace finite polyhedrality in Section 5, while a uniform barycentric interior condition replaces automatic finite-dimensional slope control.

5 The compact continuum-state prediction market

An expectation-level separation places a target quote field outside the price range of a compact set of pricing laws. Such a separation is not, by itself, a trading theorem. This section specifies the additional market object whose polar consists exactly of tradeable dynamic gains. The state and successor sets may be continua, but the trading calendar and the dynamically traded menu are finite. Continuous-time cylindrical curve trading is deliberately not treated here; under one law it is available in the literature (Kardaras, 2024; De Donno and Pratelli, 2005), and under model uncertainty it remains open.

5.1 Histories, successors, and common-numeraire prices

Fix dates

$$0 = t_0 < t_1 < \cdots < t_N = T.$$

Let $H_0 = \{h_0\}$. Recursively, let E_{n+1} be a compact metric space, let

$$\Gamma_n : H_n \rightrightarrows E_{n+1} \tag{36}$$

be nonempty, compact-valued, and continuous (both upper and lower hemicontinuous), and put

$$H_{n+1} = \text{Gr}(\Gamma_n) = \{(h, e) : h \in H_n, e \in \Gamma_n(h)\}. \tag{37}$$

Thus every H_n is compact metric. The prediction set is $\Omega = H_N$, equipped with its raw prefix filtration $\mathbb{F} = (\mathcal{F}_n)_{n=0}^N$. A history includes both the primitive price field and the progressively revealed fixing tape. In particular, $\Gamma_n(h)$ is allowed to depend only on the tape already present in h ; it cannot inspect a future fixing.

Choose one strictly positive reference numeraire and express every traded asset in its units. Let

$$S_n : H_n \longrightarrow \mathbb{R}^d, \quad n = 0, \dots, N,$$

be the resulting continuous vector of *cum-dividend* prices. It is a finite executable menu of bonds, coupon accounts, collateral-conversion accounts, or other primitive traded assets. For $(h, e) \in H_{n+1}$, write

$$\Delta S_n(h, e) = S_{n+1}(h, e) - S_n(h).$$

All deterministic affine redundancies are removed. More precisely, for each n fix a linear subspace $D_n \subseteq \mathbb{R}^d$ such that every increment lies in D_n . Positions may be projected onto D_n , because the orthogonal component has identically zero gain.

This is the same normalized primitive price vector denoted by X in [Section 2](#); the letter S is used in the Bellman proof to emphasize its role as the dynamically traded martingale.

The use of a common numeraire is load-bearing. Separate martingale measures for separate collateral conventions cannot be combined in one self-financing gain. In a term-structure setting, a known coupon is carried through its payment date by

$$G_t^{c,\kappa} = C_\kappa Z_t^c(p_\kappa), \quad \theta_\kappa \leq t \leq p_\kappa,$$

and is then stopped. The corresponding stopped, common-numeraire price must be a component of S , or its carry must be an explicitly displayed predictable gain in components of S . Counting both the stopped cum-dividend price and a separate dividend would double count carry.

Assumption 5.1 (Uniform local robust no-arbitrage). *For every $n < N$ there is $\kappa_n > 0$ such that, for every $h \in H_n$,*

$$\{z \in D_n : |z| \leq \kappa_n\} \subseteq \text{conv}\{\Delta S_n(h, e) : e \in \Gamma_n(h)\}. \tag{38}$$

The relative ball is taken in the active space D_n , not in $\mathbb{R}^d$. The assumption is stronger than mere one-step absence of arbitrage. Its uniform radius supplies attainment and a uniform bound on hedge ratios. The boundary example in [Section 8](#) shows why it cannot simply be deleted from the exact-attainment statement.

5.2 The full rectangular martingale class

We use three standard compact-state facts in their following precise forms. Berge's theorem says that the maximum of a continuous function over a nonempty compact-valued continuous correspondence is continuous. The Arsenin–Kunugui Borel uniformization theorem says that a Borel subset of a standard-Borel–Polish product whose nonempty sections are sigma-compact admits a Borel selector; the compact sections below are sigma-compact. Finally, the finite-horizon Ionescu–Tulcea theorem says that an initial law and a finite sequence of Borel probability kernels determine a unique law on the product history space. These are the only measurable-selection and kernel-concatenation results used below; see Kurka (2016, Theorem 2.1) for Arsenin–Kunugui and Kallenberg (2021) for finite kernel concatenation. We use Aliprantis and Border (2006) separately for the compact-correspondence facts.

For $h \in H_n$, define the local martingale-kernel set

$$\mathcal{P}_n(h) = \left\{ p \in \mathcal{P}(E_{n+1}) : p(\Gamma_n(h)) = 1, \int \Delta S_n(h, e) p(de) = 0 \right\}. \quad (39)$$

By (38), this set is nonempty. It is compact, and its graph is closed in $H_n \times \mathcal{P}(E_{n+1})$. Indeed, support is preserved under a joint limit by upper hemicontinuity of Γ_n , and the barycenter condition passes to the limit by continuity of ΔS_n . For the second assertion, first extend ΔS_n continuously from the closed graph $H_{n+1} \subseteq H_n \times E_{n+1}$ to the ambient product by the Tietze theorem, exactly as in the proof of Theorem 6.6; the extension is defined on a common domain, is uniformly continuous on the compact product, and therefore lets one compare $\int \Delta S_n(h^j, \cdot) dp^j$ with $\int \Delta S_n(h, \cdot) dp$. The value of the integral is unaffected, because each p^j is carried by $\Gamma_n(h^j)$, where the extension agrees with ΔS_n .

Let

$$\mathcal{M}_\Omega = \{Q \in \mathcal{P}(\Omega) : S \text{ is a } Q\text{-martingale in the raw canonical filtration}\}. \quad (40)$$

Equivalently, $Q \in \mathcal{M}_\Omega$ if and only if

$$E_Q[\varphi(h_n) \cdot \Delta S_n] = 0 \quad \text{for every } n < N \text{ and every } \varphi \in C(H_n; D_n). \quad (41)$$

Continuous tests suffice: bounded Borel tests follow by the monotone-class theorem, coordinate by coordinate.

Lemma 5.2 (Compactness, saturation, and rectangularity). *Under Assumption 5.1, $\mathcal{M}_\Omega$ is a nonempty compact convex subset of $\mathcal{P}(\Omega)$. Every sequence of Borel selectors*

$$p_n(h) \in \mathcal{P}_n(h), \quad n < N,$$

can be concatenated to give a member of $\mathcal{M}_\Omega$. Conversely, the regular conditional successor law of any $Q \in \mathcal{M}_\Omega$ belongs to $\mathcal{P}_n(h)$ for Q -almost every history h .

Proof. The graph of $\mathcal{P}_n$ is Borel with nonempty compact sections, so it has a Borel selector. Concatenating one selector at each date gives a law on Ω . Conditional mean zero at every node makes S a martingale, which proves nonemptiness and the first rectangularity assertion.

The space $\mathcal{P}(\Omega)$ is compact because Ω is compact. Condition (41) is an intersection of closed affine conditions, and hence defines a closed convex subset. This proves compactness and convexity. Finally, disintegrate any $Q \in \mathcal{M}_\Omega$ along the finite raw filtration. Its conditional successor law is supported by $\Gamma_n(h)$, and the martingale property gives the zero-barycenter restriction for almost every history. Thus it lies in $\mathcal{P}_n(h)$ almost surely. $\square$

Lemma 5.3 (Every declared path is martingale-efficient). *Under [Assumption 5.1](#), for every $\omega \in \Omega$ there is a law $Q^\omega \in \mathcal{M}_\Omega$ with $Q^\omega(\{\omega\}) > 0$. Hence the efficient scenario set generated by the full martingale class is all of Ω .*

Proof. Fix a prefix h_n of ω , let e_{n+1} be its next successor, and put $a = \Delta S_n(h_n, e_{n+1})$. Choose $\varepsilon \in (0, 1)$ small enough that

$$z = -\frac{\varepsilon}{1-\varepsilon}a \quad \text{satisfies} \quad |z| \leq \kappa_n.$$

By (38), z is the barycenter of a finitely supported probability q on $\Gamma_n(h_n)$. Therefore

$$p_n^\omega = \varepsilon \delta_{e_{n+1}} + (1-\varepsilon)q \in \mathcal{P}_n(h_n)$$

and it charges the prescribed successor. At the finitely many prefixes of ω , use these kernels; away from those prefixes use arbitrary Borel selectors of $\mathcal{P}_n$. Finite modification on singleton histories preserves Borel measurability. Concatenation gives a law in $\mathcal{M}_\Omega$, and the probability of ω is bounded below by the product of the finitely many positive ε 's. $\square$

General nondominated robust duality is developed, under substantially more flexible analytic hypotheses, by [Bouchard and Nutz \(2015\)](#). The efficient scenario qualification in pointwise arbitrage theory is emphasized by [Burzoni et al. \(2019\)](#). The compact finite-calendar theorem below uses the stronger uniform relative-interior hypothesis because it yields bounded attained slopes, continuous epsilon-selectors, and, by [Lemma 5.3](#), a certificate on every declared path rather than only on a martingale-efficient subset.

The adjective *full* in (40) is essential. The class contains every martingale law supported by the declared prediction set, not merely a preferred model library, its convex hull, or a statistically selected subset. In particular, a compact model-class envelope $\mathcal{M}$ may replace the dual class in the dynamic theorem only under the saturation identity

$$\mathcal{M} = \mathcal{M}_\Omega. \tag{42}$$

For a particular continuous separator, it is instead enough to establish its strict lower bound over $\mathcal{M}_\Omega$, either by establishing it there directly or by a separate valid transfer from a smaller class (proved equality of the relevant lower values is one such transfer). Expectation separation over a strict subset alone need not have a tradeable pathwise interpretation.

5.3 Admissible gains and carried static claims

A predictable strategy is a sequence $H = (H_n)_{n < N}$, where $H_n : \mathbf{H}_n \rightarrow D_n$ is Borel. Its normalized terminal gain is

$$(H \cdot S)_N(\omega) = \sum_{n=0}^{N-1} H_n(h_n) \cdot \Delta S_n(h_n, e_{n+1}). \tag{43}$$

All strategies constructed below are bounded. Since S is bounded on the compact histories, their gains are integrable and

$$E_Q[(H \cdot S)_N] = 0, \quad Q \in \mathcal{M}_\Omega. \tag{44}$$

Write $\mathcal{H}^b$ for the bounded Borel-predictable class and $\mathcal{H}^c \subseteq \mathcal{H}^b$ for the subclass having continuous H_n .

Every quoted static object used in the Bellman and separation argument must have a continuous common-unit representative $g \in C(\Omega)$. This representative may be the actual terminal ledger when

that ledger is continuous. If a physically delivered ledger is discontinuous because of its exercise indicator, use instead its continuous fixing-date mark g^{fix} ; the market bid and ask remain those of the actual contract. In either case the required carry interface is an exact pathwise identity

$$g^{\text{pay}} = g^{\text{fix}} + (K^g \cdot S)_N \quad \text{on } \Omega \quad (45)$$

for a declared bounded predictable strategy K^g , after both sides have been put in the same numeraire. Thus a subhedge of g^{fix} becomes a subhedge of g^{pay} by adding K^g to the dynamic strategy, without changing the quoted static cost. It is not enough that the two expressions have the same expectation under a preferred model library.

For a finite piecewise-linear potential, the canonical decomposition is (34). It is executable only when the resulting weighted cash, forward, and call or put components are listed instruments with their own bid and ask. In particular, the constant component cannot be dropped on a variable-mass stratum, and the cost is the sign-dependent cost of the primitive positions rather than a fictitious mid-price integral of the potential.

6 Pointwise robust duality and executable infeasibility

This section is the continuum counterpart of Section 3. The prediction set is now the compact history space of Section 5 rather than a finite tree, so finite polyhedrality is replaced by a one-step affine duality with attained slopes, iterated by a Bellman recursion into the finite-calendar robust duality with all four extrema attained and a pathwise optional decomposition; the closed gain cone and its exact polar follow, and the section ends by reducing the infeasibility of an uncountable quote book to a finite executable certificate. The uniform relative-interior condition of assumption 5.1 is the hypothesis carrying every step.

6.1 The one-step affine dual

The following finite-dimensional lemma is the load-bearing separation result. It is stated for the upper and lower problems because the latter is the form used to turn a positive expectation gap into a trading certificate.

Lemma 6.1 (One-step martingale duality with attained slopes). *Let K be compact metric, let D be a finite-dimensional Euclidean space, and let $a : K \rightarrow D$ and $u : K \rightarrow \mathbb{R}$ be continuous. Suppose that*

$$\{z \in D : |z| \leq \kappa\} \subseteq \text{conv } a(K) \quad \text{for some } \kappa > 0.$$

Put

$$\mathcal{P}_0 = \left\{ p \in \mathcal{P}(K) : \int a \, dp = 0 \right\}.$$

Then all four extrema below are attained and

$$\min_{H \in D} \max_{y \in K} \{u(y) - H \cdot a(y)\} = \max_{p \in \mathcal{P}_0} \int u \, dp, \quad (46)$$

$$\max_{H \in D} \min_{y \in K} \{u(y) - H \cdot a(y)\} = \min_{p \in \mathcal{P}_0} \int u \, dp. \quad (47)$$

An optimizing slope may be chosen with

$$|H| \leq \frac{2\|u\|_\infty}{\kappa}. \quad (48)$$

Proof. The set $\mathcal{P}_0$ is nonempty and compact. Consider

$$C = \left\{ \left(\int a \, dp, \int u \, dp \right) : p \in \mathcal{P}(K) \right\} \subseteq D \times \mathbb{R}.$$

It is compact and convex. For $x \in \text{conv } a(K)$, define the upper boundary

$$\hat{u}(x) = \max \{ r : (x, r) \in C \}.$$

This is a finite upper-semicontinuous concave function. Since zero lies in the relative interior of its domain, the supporting-hyperplane theorem for the hypograph gives a supergradient $H \in D$ at zero:

$$\hat{u}(x) \leq \hat{u}(0) + H \cdot x.$$

Apply this inequality to the Dirac measure at y . With $m = \hat{u}(0) = \max_{p \in \mathcal{P}_0} \int u \, dp$, it gives

$$u(y) - H \cdot a(y) \leq m \quad (y \in K).$$

Thus the left side of (46) is at most m . The reverse inequality follows by integrating $u - H \cdot a$ against any $p \in \mathcal{P}_0$. Compactness gives attainment of the measure maximum.

Let $F(H) = \max_y \{u(y) - H \cdot a(y)\}$. Since the convex hull of $a(K)$ contains the radius- κ ball,

$$F(H) \geq -\|u\|_\infty + \kappa|H|, \quad F(0) \leq \|u\|_\infty.$$

Consequently a minimizer exists in the ball displayed in (48). Applying the upper identity to $-u$, then changing the sign of H , proves (47), its attainment, and the same bound. $\square$

6.2 Two-sided slope estimates and the scope of the radius

The bound (48) is a position limit: no one-step robust hedge needs a holding larger than twice the payoff's size divided by the smallest guaranteed move. We now ask how sharp it is, and in what sense the radius is the right quantity. Keep the setting of Lemma 6.1 and first suppose that $\dim D \geq 1$. Write $C = \text{conv } a(K)$, and let

$$h(v) := \max_{y \in K} v \cdot a(y), \quad \rho := \min_{|v|=1} h(v), \quad \bar{\rho} := \min_{|v|=1} \max\{h(v), h(-v)\},$$

be the support function of C , the largest radius of a ball around 0 contained in C (so $\rho \geq \kappa$), and its two-sided counterpart; clearly $\rho \leq \bar{\rho}$. For $u \in C(K)$ let $V(u)$ denote the common value in (46), $\mathcal{O}(u) := \{H \in D : u(y) - H \cdot a(y) \leq V(u) \, \forall y \in K\}$ the set of optimal slopes — nonempty, closed, and convex — and $H^\circ(u)$ its element of least norm. The *slope constant* of the geometry is

$$\Lambda := \sup\{|H^\circ(u)| : u \in C(K), \|u\|_\infty \leq 1\}.$$

If $D = \{0\}$, then $a \equiv 0$, $V(u) = \max_K u$, $\mathcal{O}(u) = \{0\}$, and $\Lambda = 0$; every finite-radius ball in D is just $\{0\}$. Thus zero-dimensional dates have identically zero holdings and are omitted from the unit-sphere radius ratios below.

Proposition 6.2 (Two-sided slope estimate). *1. For every $u \in C(K)$ and every $H \in \mathcal{O}(u) \setminus \{0\}$, with $v := H/|H|$,*

$$|H| \leq \frac{V(u) - \min_K u}{h(-v)} \leq \frac{\text{osc } u}{h(-v)} \leq \frac{\text{osc } u}{\rho} \leq \frac{2\|u\|_\infty}{\rho}.$$

2. For every unit $v \in D$, the payoff $u_v := (v \cdot a) / \max\{h(v), h(-v)\}$ satisfies $\|u_v\|_\infty = 1$, $V(u_v) = 0$, and $\mathcal{O}(u_v) = \{v / \max\{h(v), h(-v)\}\}$. Consequently $\Lambda \geq 1/\bar{\rho}$.
3. Hence $1/\bar{\rho} \leq \Lambda \leq 2/\rho$. If C is centrally symmetric, then $\bar{\rho} = \rho$ and $1/\rho \leq \Lambda \leq 2/\rho$: the inscribed radius determines the slope constant up to a factor of two.

Proof. (i) By compactness choose $y^- \in K$ with $-v \cdot a(y^-) = h(-v)$; since $B_\rho \subseteq C$, $h(-v) \geq \rho > 0$. The constraint at y^- reads $u(y^-) + |H| h(-v) \leq V(u)$, whence $|H| h(-v) \leq V(u) - u(y^-) \leq V(u) - \min_K u$. Taking $H = 0$ in the primal gives $V(u) \leq \max_K u$, so $|H| h(-v) \leq \text{osc } u$.

(ii) Put $c := 1/\max\{h(v), h(-v)\}$, so $\|u_v\|_\infty = 1$. For any $H \in D$, $\max_y \{u_v(y) - H \cdot a(y)\} = h(cv - H)$, the support function of C at $cv - H$. Because $B_\rho \subseteq C$, $h(w) \geq \rho|w|$ with equality to zero only at $w = 0$; hence the primal minimum is 0, attained exactly at $H = cv$, and $V(u_v) = 0$ (the dual is $\max_{p \in \mathcal{P}_0} c \int v \cdot a \, dp = 0$). Thus $H^\circ(u_v) = cv$ and $|H^\circ(u_v)| = c$; minimizing over v gives $\Lambda \geq 1/\bar{\rho}$.

(iii) Combine (i) with u ranging over the unit ball and (ii). For centrally symmetric C , $h(v) = h(-v)$ for all v , so $\bar{\rho} = \rho$. $\square$

Example 6.3 (The radius is sufficient, not necessary, for bounded slopes). Let $K = \{y_1, y_2\}$, $D = \mathbb{R}$, $a_\varepsilon(y_1) = -1$, and $a_\varepsilon(y_2) = \varepsilon$. For $\varepsilon \in (0, 1]$, $\rho_\varepsilon = \varepsilon$ while $\bar{\rho}_\varepsilon = 1$. For $u = (u_1, u_2)$ the unique zero-mean kernel has weights $\varepsilon/(1 + \varepsilon)$ and $1/(1 + \varepsilon)$, so $V_\varepsilon(u) = (\varepsilon u_1 + u_2)/(1 + \varepsilon)$, and the two constraints $u_1 + H \leq V_\varepsilon(u)$, $u_2 - \varepsilon H \leq V_\varepsilon(u)$ force $\mathcal{O}_\varepsilon(u) = \{(u_2 - u_1)/(1 + \varepsilon)\}$. Hence $\Lambda_\varepsilon = 2/(1 + \varepsilon) \leq 2$.

The family extends continuously to $\varepsilon = 0$. There $V_0(u) = u_2$, $\mathcal{O}_0(u) = (-\infty, u_2 - u_1]$, and $H_0^\circ(u) = \min\{0, u_2 - u_1\}$, so $\rho_0 = 0$, $\bar{\rho}_0 = 1$, and $\Lambda_0 = 2$. Thus the inscribed radius can collapse while every least-norm attaining slope stays bounded by 2. For $\varepsilon > 0$ the geometry is increasingly unbalanced, $h(1)/h(-1) = \varepsilon$, and [Proposition 6.2\(ii\)](#) only guarantees $\Lambda_\varepsilon \geq 1/\bar{\rho}_\varepsilon = 1$.

For a genuinely uniform statement, let I be any nonempty index set of one-step geometries (K_i, D_i, a_i) satisfying the hypotheses above, with $\dim D_i \geq 1$ and zero in the D_i -interior of $\text{conv } a_i(K_i)$. Attach to geometry i the quantities $h_i, \rho_i, \bar{\rho}_i, \Lambda_i$, and put

$$\rho_* := \inf_{i \in I} \rho_i, \quad \Lambda_* := \sup_{i \in I} \Lambda_i.$$

For $M \geq 1$, call the family *uniformly M -balanced* if $h_i(v) \leq M h_i(-v)$ for every $i \in I$ and every unit $v \in D_i$: opposite moves are available in comparable size with one common constant.

Theorem 6.4 (Uniform slope–radius estimates under balance). 1. If $\rho_* > 0$, then $\Lambda_* \leq 2/\rho_*$. In particular, a common positive radius bounds every least-norm attaining one-step slope by $2\|u\|_\infty/\rho_*$, uniformly over the family.

2. If the family is uniformly M -balanced and $\Lambda_* < \infty$, then

$$\rho_* \geq \frac{1}{M\Lambda_*} > 0.$$

Thus, under a common balance bound, uniformly bounded least-norm optimizers imply a common positive radius.

3. Without a common balance bound the converse in (ii) fails even when every individual geometry has zero in its interior: take the subfamily of [Example 6.3](#) with $I = \mathbb{N}$ and $\varepsilon_i = 1/i$. Then $\rho_* = 0$ while $\Lambda_* = 2$.

Proof. (i) is [Proposition 6.2](#)(i), applied to every i . For (ii), uniform M -balance gives, for every unit $v \in D_i$,

$$\max\{h_i(v), h_i(-v)\} \leq M \min\{h_i(v), h_i(-v)\},$$

and hence $\bar{\rho}_i \leq M\rho_i$. By [Proposition 6.2](#)(ii), $\Lambda_i \geq 1/\bar{\rho}_i \geq 1/(M\rho_i)$, so $\rho_i \geq 1/(M\Lambda_i) \geq 1/(M\Lambda_*)$. Take the infimum over i . Part (iii) is the stated subfamily of the example; its smallest balance constants diverge. $\square$

Remark 6.5 (What the radius is, and is not). At a fixed market date with $d_n \geq 1$, apply the theorem to the family indexed by $h \in \mathbf{H}_n$. Assumption (38) gives $\rho_* \geq \kappa_n$. More generally, under the compact-continuous hypotheses of [Section 5](#),

$$(h, v) \mapsto \max_{e \in \Gamma_n(h)} v \cdot \Delta S_n(h, e)$$

is continuous by Berge's theorem. Therefore $h \mapsto \rho_n(h)$ is continuous, and pointwise interiority on compact $\mathbf{H}_n$ already forces a positive uniform minimum. The positive- ε family in part (iii) consequently cannot be embedded as histories of one such market without acquiring a boundary limit. The full $\varepsilon \in [0, 1]$ family of [Example 6.3](#) does embed compact-continuously: take two dates, $\Gamma_0(h_0) = [0, 1]$, $S_0 = S_1 \equiv 0$, $\Gamma_1(\varepsilon) = \{y_1, y_2\}$, and $S_2(\varepsilon, y_1) = -1$, $S_2(\varepsilon, y_2) = \varepsilon$. Its endpoint has $\rho_0 = 0$ and lies exactly outside [Assumption 5.1](#); nevertheless its least-norm optimizers remain uniformly bounded.

Thus the uniform radius is the clean sufficient condition, and across arbitrary pointwise-interior families a common balance constant makes it quantitatively equivalent to uniformly bounded least-norm optimizers. Without uniform balance it is not necessary; the two-sided $\bar{\rho}$ supplies the universal lower estimate and records the asymmetry. The defective nodes of [Proposition 7.3](#) are symmetric, $\{\pm\varepsilon_j\}$, so there radius collapse and slope blow-up are equivalent; the augmentation of [Lemma 7.4](#) restores a uniform radius, which by (i) is all that the convergence theory needs, without asserting balance.

6.3 Bellman duality and robust optional decomposition

For $f \in C(\Omega)$, define recursively

$$\bar{V}_N = f, \quad \bar{V}_n(h) = \max_{p \in \mathcal{P}_n(h)} \int \bar{V}_{n+1}(h, e) p(\mathrm{d}e), \quad (49)$$

$$\underline{V}_N = f, \quad \underline{V}_n(h) = \min_{p \in \mathcal{P}_n(h)} \int \underline{V}_{n+1}(h, e) p(\mathrm{d}e). \quad (50)$$

Theorem 6.6 (Compact finite-calendar robust duality). *Under the hypotheses of [Section 5](#), for every $f \in C(\Omega)$,*

$$\bar{V}_0(h_0) = \max_{Q \in \mathcal{M}_\Omega} E_Q f = \min \left\{ x : \exists H \in \mathcal{H}^b, \quad x + (H \cdot S)_N \geq f \text{ on } \Omega \right\}, \quad (51)$$

$$\underline{V}_0(h_0) = \min_{Q \in \mathcal{M}_\Omega} E_Q f = \max \left\{ x : \exists H \in \mathcal{H}^b, \quad x + (H \cdot S)_N \leq f \text{ on } \Omega \right\}. \quad (52)$$

Both model extrema and both hedging extrema are attained. Every Bellman function in (49)–(50) is continuous.

For the upper problem there are a bounded Borel-predictable H and an adapted nondecreasing process K , with $K_0 = 0$, such that, pathwise,

$$\bar{V}_n = \bar{V}_0 + (H \cdot S)_n - K_n, \quad n = 0, \dots, N. \quad (53)$$

Moreover, for every $\epsilon > 0$, the upper and lower inequalities admit continuous predictable strategies after worsening the initial value by at most ϵ . Thus strict certificates can always be made continuous.

Proof. We prove the upper statement; the lower statement follows by applying it to $-f$. Apply [Lemma 6.1](#) at a history h , with $K = \Gamma_n(h)$, $a(e) = \Delta S_n(h, e)$, and $u(e) = \bar{V}_{n+1}(h, e)$. It gives the equivalent recursion

$$\bar{V}_n(h) = \min_{H \in D_n} \max_{e \in \Gamma_n(h)} \{\bar{V}_{n+1}(h, e) - H \cdot \Delta S_n(h, e)\}. \quad (54)$$

Inductively, $\|\bar{V}_n\|_\infty \leq \|f\|_\infty$. The slope bound permits restriction of the minimum to the fixed compact ball

$$\{H \in D_n : |H| \leq 2\|f\|_\infty/\kappa_n\}.$$

If $\bar{V}_{n+1}$ is continuous, Berge's maximum theorem and the continuity of Γ_n show that the inner maximum in (54) is jointly continuous in (h, H) . Minimizing over the fixed compact ball proves that $\bar{V}_n$ is continuous. Backward induction starts from f .

Extend the continuous function on the closed graph $H_{n+1} \subseteq H_n \times E_{n+1}$ continuously to the ambient product. It follows that integration against a weakly varying kernel is jointly continuous on the graph of $\mathcal{P}_n$. The argmax graph in the measure form (49) is therefore Borel with nonempty compact sections, so the measurable maximum theorem supplies a Borel maximizing kernel. Concatenate these kernels. The resulting law Q^* belongs to $\mathcal{M}_\Omega$ by [Lemma 5.2](#), and iterated conditioning gives $E_{Q^*}f = \bar{V}_0$. For any other $Q \in \mathcal{M}_\Omega$, its conditional kernels lie in $\mathcal{P}_n$ almost surely, so backward conditioning gives $E_Q f \leq \bar{V}_0$. This proves the model equality and its attainment.

The argmin graph in (54) is likewise Borel with nonempty compact sections. Choose a bounded Borel selector $H_n(h)$. For every allowed successor,

$$\bar{V}_n(h) + H_n(h) \cdot \Delta S_n(h, e) \geq \bar{V}_{n+1}(h, e). \quad (55)$$

Summing proves the pointwise superhedge with initial capital $\bar{V}_0$. Conversely, take expectations of any pointwise superhedge under any $Q \in \mathcal{M}_\Omega$ and use (44). Its initial capital is at least $\max_Q E_Q f = \bar{V}_0$. This proves the hedging equality and attainment.

Define

$$K_{n+1} - K_n = \bar{V}_n + H_n \cdot \Delta S_n - \bar{V}_{n+1}.$$

It is nonnegative at every path by (55), and rearrangement and summation give (53).

It remains to prove the continuous-strategy assertion. For a fixed date, write the continuous convex objective in (54) as $F_n(h, H)$. At every h , choose an optimizer H^h . Continuity of F_n and $\bar{V}_n$ gives a neighborhood U_h on which

$$F_n(h', H^h) \leq \bar{V}_n(h') + \epsilon/N.$$

Take a finite subcover and a continuous partition of unity $(\rho_j)_j$ subordinate to it. The convex combination $H_n^c(h') = \sum_j \rho_j(h') H^{h_j}$ is continuous and, by convexity in the slope,

$$F_n(h', H_n^c(h')) \leq \sum_j \rho_j(h') F_n(h', H^{h_j}) \leq \bar{V}_n(h') + \epsilon/N.$$

Summing these N approximate node inequalities gives $\bar{V}_0 + \epsilon + (H^c \cdot S)_N \geq f$ on all of Ω . Applying the same argument to $-f$ proves the lower version. $\square$

6.4 The closed gain cone

The theorem proves the closedness needed here without asserting convergence of the strategies themselves. Define

$$\mathcal{C}_0 = \{f \in C(\Omega) : \exists \text{ bounded Borel-predictable } H, (H \cdot S)_N \geq f \text{ on } \Omega\}. \quad (56)$$

Corollary 6.7 (Closed-cone and polar identities). *One has*

$$\mathcal{C}_0 = \{f \in C(\Omega) : \sup_{Q \in \mathcal{M}_\Omega} E_Q f \leq 0\}. \quad (57)$$

Consequently $\mathcal{C}_0$ is a norm-closed convex cone. In addition,

$$\mathcal{C}_0 = \overline{\{(H \cdot S)_N - k : H \in \mathcal{H}^c, k \in C(\Omega), k \geq 0\}}^{\|\cdot\|_\infty}. \quad (58)$$

The probability measures in the normalized polar of this cone are exactly $\mathcal{M}_\Omega$.

Proof. If $f \in \mathcal{C}_0$, expectation under every martingale law is nonpositive. Conversely, if the right side of (57) holds, Theorem 6.6 gives $x + (H \cdot S)_N \geq f$ with $x \leq 0$; hence the same strategy with zero initial cash also dominates f . This proves (57), and norm closedness follows because $f \mapsto \sup_Q E_Q f$ is 1-Lipschitz in the uniform norm.

The inclusion “ $\supseteq$ ” in (58) follows from martingale expectations and closedness. For the reverse inclusion, take $f \in \mathcal{C}_0$. The continuous-strategy part of Theorem 6.6 gives $\epsilon + (H^\epsilon \cdot S)_N \geq f$. Thus

$$f - \epsilon = (H^\epsilon \cdot S)_N - k_\epsilon, \quad k_\epsilon = \epsilon + (H^\epsilon \cdot S)_N - f \geq 0,$$

and k_ϵ is continuous. Let $\epsilon \downarrow 0$.

Finally, a normalized positive functional annihilates every continuous gain if and only if it satisfies the continuous martingale tests (41). By the Riesz representation theorem and Lemma 5.2, the normalized polar is exactly $\mathcal{M}_\Omega$. $\square$

6.5 From continuum quote infeasibility to a finite trade

Let $\mathfrak{A}$ be any index set of actually quoted contracts, each represented in the robust dual by a continuous normalized mark $g_a \in C(\Omega)$. When the actual payment ledger differs from this mark, assume the exact carry identity (45); the bid and ask below are the prices of the actual contract. They satisfy

$$-\infty < \underline{q}_a \leq \bar{q}_a < \infty.$$

Define

$$\mathcal{Q}_{\text{cal}} = \left\{ Q \in \mathcal{M}_\Omega : \underline{q}_a \leq E_Q g_a \leq \bar{q}_a \text{ for every } a \in \mathfrak{A} \right\}. \quad (59)$$

For a finite list $J = (a_1, \dots, a_q)$ and $\theta \in \mathbb{R}^q$, its executable bid–ask cost is

$$c_J(\theta) = \sum_{i=1}^q \left(\theta_i^+ \bar{q}_{a_i} - \theta_i^- \underline{q}_{a_i} \right) = \sup_{y \in \prod_i [\underline{q}_{a_i}, \bar{q}_{a_i}]} \theta \cdot y. \quad (60)$$

Theorem 6.8 (Finite executable certificate for a continuum quote book). *Under the hypotheses of Theorem 6.6, the following are equivalent.*

- (i) $\mathcal{Q}_{\text{cal}} = \emptyset$.

- (ii) There are a finite list $J = (a_1, \dots, a_q)$, a vector $\theta \in \mathbb{R}^q$, a bounded Borel-predictable strategy H , and numbers $x \in \mathbb{R}$, $\delta > 0$ such that

$$x + c_J(\theta) < 0, \quad (61)$$

$$x + \sum_{i=1}^q \theta_i g_{a_i} + (H \cdot S)_N \geq \delta \quad \text{on every } \omega \in \Omega. \quad (62)$$

For the continuous representative marks, the strategy in (ii) may be chosen continuous. Thus infeasibility of even an uncountable quote family has a finite, pointwise, semi-static and dynamic certificate. Equality quotes are included by setting $\underline{q}_a = \bar{q}_a$. If a selected g_{a_i} is a fixing-mark representative, replacing it by its actual payment ledger and subtracting $\theta_i K^{a_i}$ from the dynamic strategy preserves the initial cost and the terminal margin. The resulting actual-ledger strategy is bounded and Borel; it is continuous when the declared carry strategies are continuous.

Proof. For each a , the set

$$F_a = \{Q \in \mathcal{M}_\Omega : \underline{q}_a \leq E_Q g_a \leq \bar{q}_a\}$$

is closed in the compact set $\mathcal{M}_\Omega$. If their total intersection is empty, the finite-intersection property supplies $a_1, \dots, a_q$ whose intersection is already empty.

Let

$$A_J = \{(E_Q g_{a_i})_{i=1}^q : Q \in \mathcal{M}_\Omega\}, \quad B_J = \prod_{i=1}^q [\underline{q}_{a_i}, \bar{q}_{a_i}].$$

These are disjoint compact convex subsets of $\mathbb{R}^q$. Strict separation, with orientation changed if necessary, gives θ and

$$\gamma := \min_{z \in A_J} \theta \cdot z - \max_{y \in B_J} \theta \cdot y > 0. \quad (63)$$

In view of (60), the continuous payoff

$$f_\theta = \sum_i \theta_i g_{a_i} - c_J(\theta)$$

satisfies $\min_{Q \in \mathcal{M}_\Omega} E_Q f_\theta = \gamma$. The subhedging identity (52), followed by a change of sign of the selected trading strategy, supplies a bounded Borel strategy H such that

$$f_\theta + (H \cdot S)_N \geq \gamma \quad \text{on } \Omega. \quad (64)$$

Choose $x = -c_J(\theta) - \gamma/2$. Then the initial cost is $-\gamma/2$, while the terminal payoff is at least $\gamma/2$. This is (ii). To obtain a continuous strategy, use the continuous $\gamma/2$ -approximate subhedge and choose $x = -c_J(\theta) - \gamma/4$; both the initial receipt and the terminal margin are then at least $\gamma/4$ in absolute value with the desired sign.

For a representative fixing mark, multiply its carry identity by θ_i and sum. The actual-ledger payoff is obtained by subtracting $\sum_i \theta_i K^{a_i}$ from the strategy in (62). This is bounded and predictable, and the pathwise identity leaves both cost and margin unchanged.

Conversely, suppose $Q \in \mathcal{Q}_{\text{cal}}$. The static expectation satisfies

$$E_Q \sum_i \theta_i g_{a_i} \leq c_J(\theta),$$

and the dynamic gain has expectation zero. A portfolio with negative initial cost and a strictly positive pointwise terminal payoff contradicts these two facts. Hence (ii) implies (i). $\square$

Remark 6.9 (Exact interface with an expectation-level separation). The theorem applies directly to the original finite call, caplet, swaption, mass, and forward quote coordinates. A bounded-Lipschitz potential may be substituted only after (34) and (45) have converted it into a finite portfolio of listed claims and an admissible carry gain. Separation over a preferred model envelope must also imply separation over the saturated class $\mathcal{M}_\Omega$, either directly or through a separate valid transfer (for example, proved equality of the relevant robust values), as required by (42). Under those two interfaces, the strictly positive finite-level separation gap is the γ in (63), and Theorem 6.8 completes it to a trade.

7 Finite-tree limits and faithful discretization

The finite-tree linear programmes of Section 3 and the continuum-state theorem of Section 6 live on different state spaces, and this section builds the bridge between them. Weak density of a preferred model library, even when proved, only identifies expectations of continuous test payoffs; it neither constructs nested finite trees whose full martingale classes recover the continuum class nor controls bounded strategies that remain valid off the grid. The bridge built here proves the stronger law-set recovery first and then restores the local hedging geometry.

7.1 Martingale cubature and recovery of the full law class

We first prove the bridge rather than assume it. The proof is a raw-history version of martingale Tchakaloff cubature (for the classical theorem see Tchakaloff, 1957; Bayer and Teichmann, 2006; for its martingale form, Beiglböck and Nutz, 2014, Theorem 5.1; for the cubature paradigm on path space, Lyons and Victoir, 2004): the atoms are selected jointly with the progressively revealed tape, so no future fixing is inserted into an earlier sigma-field.

Lemma 7.1 (Finite raw-history martingale cubature). *Let $Q \in \mathcal{M}_\Omega$ and let $f^1, \dots, f^r \in C(\Omega)$. There are a finite rooted subtree $\Omega^{Q,f} \subseteq \Omega$ and a law $Q^f \in \mathcal{M}_\Omega$, supported on that subtree, such that*

$$E_{Q^f} f^j = E_Q f^j, \quad j = 1, \dots, r. \quad (65)$$

At each selected node the successors can be taken from a prescribed full-conditional-measure set. In particular, all pathwise structural and tape identities defining Ω are preserved.

Proof. Because the history spaces are compact metric, choose Borel regular conditional successor kernels p_n of Q , and choose bounded Borel versions

$$v_n^j(h) = E_Q[f^j \mid \mathcal{F}_n](h), \quad v_N^j = f^j.$$

They satisfy $|v_n^j| \leq \|f^j\|_\infty$. Let Q_n be the history marginal at time n . Set $\mathbf{G}_N = \mathbf{H}_N$. Recursively backward, let $\mathbf{G}_n$ be the set of histories h at which $p_n(h)$ is supported by $\Gamma_n(h)$, has zero ΔS_n -barycenter, satisfies all tower identities

$$\int v_{n+1}^j(h, e) p_n(h, de) = v_n^j(h), \quad j = 1, \dots, r. \quad (66)$$

and obeys

$$p_n(h, \{e : (h, e) \in \mathbf{G}_{n+1}\}) = 1. \quad (67)$$

The kernel-integral maps in these conditions are Borel and the successor graphs are closed, hence every $\mathbf{G}_n$ is Borel. The martingale and tower properties, followed by backward disintegration, give

$Q_n(G_n) = 1$ at every date. In particular the deterministic root is good, and the conditional kernel at every good history is concentrated on good successors.

We use the following elementary finite-dimensional barycentric fact. If Z is a bounded Borel $\mathbb{R}^k$ -valued random vector and G is a measurable set of full probability, then $E[Z]$ is a convex combination of at most $k + 1$ points of $Z(G)$. To see this without a support qualification, intersect G with the inverse image of the essential range; this remains a full-measure set, and its Z -image is dense in the essential range. Put $C = \text{conv } Z(G)$, using this smaller full-measure set. The expectation lies in $\text{ri } \overline{C}$: otherwise a supporting functional at a proper relative boundary face would be bounded above by its value at $E[Z]$, and equality in expectation would put the essential range in that face, contradicting properness relative to its affine hull. For a nonempty finite-dimensional convex set, $\text{ri } \overline{C} = \text{ri } C \subseteq C$. Caratheodory's theorem now represents $E[Z]$ by at most $k + 1$ actual points of $Z(G)$, as claimed.

Start at h_0 . Apply this fact under $p_0(h_0, \cdot)$ to

$$Z_0(e) = (\Delta S_0(h_0, e), v_1^1(h_0, e), \dots, v_1^r(h_0, e)), \quad (68)$$

choosing atoms only in the good full-measure set. The selected weights have zero price-increment barycenter and reproduce $(v_0^1, \dots, v_0^r)$. At each of the finitely many selected successors, repeat the construction with $(\Delta S_1, v_2^1, \dots, v_2^r)$, and continue until time N . The resulting tree is finite. Its nodewise zero barycenters make the concatenated law a martingale law, and (66) telescopes to (65). Every atom was selected from an allowed successor set, so its paths remain in Ω . $\square$

This is the finite-horizon mechanism behind the martingale Tchakaloff theorem; compare [Beiglböck and Nutz \(2014, Theorem 5.1\)](#). The preceding proof records the extra raw-tape step needed here instead of applying ordinary cubature to the terminal price vector alone.

Fix a dense sequence $(\varphi_j)_{j \geq 1}$ in the unit ball of $C(\Omega)$, and use the compatible weak metric

$$d_w(P, Q) = \sum_{j \geq 1} 2^{-j} |E_P \varphi_j - E_Q \varphi_j|. \quad (69)$$

Proposition 7.2 (Nested finite trees recover the full martingale class). *There are nested finite subtrees $\Omega_m \subseteq \Omega$ such that, if $\mathcal{M}_m$ is the full set of martingale laws supported on Ω_m , then $\mathcal{M}_m \neq \emptyset$, $\mathcal{M}_m \subseteq \mathcal{M}_\Omega$, and*

$$d_H^w(\mathcal{M}_m, \mathcal{M}_\Omega) \longrightarrow 0, \quad (70)$$

where d_H^w is Hausdorff distance induced by (69).

Proof. The compact set $\mathcal{M}_\Omega$ has a finite $1/m$ -net $Q^{m,1}, \dots, Q^{m,k_m}$. Choose $L_m \uparrow \infty$ so that $2^{1-L_m} < 1/m$. Apply [Lemma 7.1](#) to each net law and the test vector $(\varphi_1, \dots, \varphi_{L_m})$. The resulting finite-tree law $\hat{Q}^{m,k}$ matches the first L_m integrals, and hence

$$d_w(Q^{m,k}, \hat{Q}^{m,k}) \leq \sum_{j > L_m} 2^{-j} 2 = 2^{1-L_m} < 1/m. \quad (71)$$

Take all prefixes and precisely the observed parent–successor edges in the supports of these cubature laws, cumulatively over stages $1, \dots, m$. They generate a nested finite subtree Ω_m . No rectangularization adds an inadmissible path: a node records its complete history, and every retained outgoing edge belongs to Γ_n at that exact history. Each $\hat{Q}^{m,k}$ remains supported on Ω_m . Every full martingale law on this subtree is supported on Ω , so $\mathcal{M}_m \subseteq \mathcal{M}_\Omega$.

Given $Q \in \mathcal{M}_\Omega$, choose a net point $Q^{m,k}$ within $1/m$. Its cubature law lies in $\mathcal{M}_m$ and is within another $1/m$ by (71). Thus

$$\sup_{Q \in \mathcal{M}_\Omega} d_w(Q, \mathcal{M}_m) \leq 2/m, \quad \sup_{P \in \mathcal{M}_m} d_w(P, \mathcal{M}_\Omega) = 0,$$

which proves (70). $\square$

7.2 Law recovery is not hedge recovery

Recovering the martingale class, and with it every robust value, does not recover the geometry that bounds the hedges. The following proposition makes this precise: a cubature sequence can satisfy (70) while the local increment hulls of its trees lose their interior radius, so that the nodewise Bellman-attaining slopes are unbounded in m .

Proposition 7.3 (Cubature can recover the martingale class while collapsing the radius). *There is a two-step compact market satisfying (38) with $\kappa_0 = \kappa_1 = 1$, together with nested finite trees $\Omega_m \subseteq \Omega$ obtainable by the cumulative-support cubature scheme of Proposition 7.2, such that the full martingale classes recover the continuum class,*

$$\sup_{Q \in \mathcal{M}_\Omega} d_w(Q, \mathcal{M}_m) \longrightarrow 0,$$

while at some second-step histories of Ω_m the local increment hull has radius $\varepsilon_j = 2^{-2j} \rightarrow 0$, and for one fixed continuous payoff the nodewise Bellman-attaining slopes on Ω_m are unbounded in m . Adjoining two successors at each defective node restores radius one and bounds those slopes uniformly.

Proof. Take two trading steps. At the first step the raw history is $h_1 = (x, u) \in [-1, 1] \times [0, 1]$, where u is a revealed but untraded tape coordinate, and set

$$S_0 = 0, \quad S_1(x, u) = x.$$

At the second step let $y \in [-1, 1]$ and set $S_2(x, u, y) = x + y$. Thus $D_0 = D_1 = \mathbb{R}$, both continuum increment hulls are $[-1, 1]$, and one may take $\kappa_0 = \kappa_1 = 1$.

Let $\delta_m = 2^{-m}$, and let the good core of Ω_m contain the dyadic δ_m -grids for x, u , with the full dyadic δ_m -grid for y after every good history. In addition, for $j \leq m$ retain the distinct raw histories

$$b_j = (0, u_j), \quad u_j = \frac{\sqrt{2}}{j+2},$$

and at b_j retain only the successors $y = \pm \varepsilon_j$, where $\varepsilon_j = 2^{-2j}$. The irrational tape values ensure that no later good grid adds an edge at a bad history. The trees are finite and nested, and their full martingale classes $\mathcal{M}_m$ are nonempty subsets of $\mathcal{M}_\Omega$.

They still recover the full continuum class. Indeed, for any $Q \in \mathcal{M}_\Omega$, let u_m be a nearest-grid rounding of u . Conditional on (x, u, y) , independently barycentrically round x and y to adjacent dyadic grid points, obtaining x_m and y_m . Then $E[x_m | x] = x$. Moreover, if a is any bounded function of the rounded first-step history, conditional independence gives

$$E[a(x_m, u_m)y_m] = E[E[a(x_m, u_m) | x, u] y] = E[E[a(x_m, u_m) | x, u] E[y | x, u]] = 0,$$

because $E[y | x, u] = 0$. Hence $E[y_m | x_m, u_m] = 0$, and the law Q_m of (x_m, u_m, y_m) belongs to $\mathcal{M}_m$. In the product metric $\rho((x, u, y), (x', u', y')) = |x - x'| + |u - u'| + |y - y'|$, this coupling

satisfies $\rho((x, u, y), (x_m, u_m, y_m)) \leq 3\delta_m$ almost surely. Uniform continuity of the functions defining d_w , followed by control of its geometric tail, therefore gives

$$\sup_{Q \in \mathcal{M}_\Omega} d_w(Q, \mathcal{M}_m) \longrightarrow 0.$$

The reverse distance is zero because $\mathcal{M}_m \subseteq \mathcal{M}_\Omega$.

This sequence is also obtainable by the same cumulative-support cubature scheme as [Proposition 7.2](#). Take finite nets from the good-grid martingale laws. Include finitely many edge-witness laws: to charge a prescribed nonzero root value x , pair it with the opposite endpoint of $[-1, 1]$ using the zero-mean weights, and use the point mass at zero when $x = 0$; at the target history, treat a prescribed y in the same way and use the zero increment at the auxiliary history. Thus every good-grid edge occurs in a finite martingale law. Also include, at stage j , the law concentrated at b_j with equal weights on $\pm \varepsilon_j$. These laws are already finitely supported, so each may be used as its own exact raw-history cubature, and their cumulative supports are precisely the trees just described.

Nevertheless the local increment hull at b_j has radius only ε_j . For the fixed continuous payoff

$$\psi(x, u, y) = \text{sgn}(y)\sqrt{|y|}, \quad \psi(x, u, 0) = 0,$$

the only zero-mean kernel on $\{\pm \varepsilon_j\}$ gives value zero. An affine subhedge attaining that nodewise value must satisfy

$$H_j \varepsilon_j \leq \sqrt{\varepsilon_j}, \quad -H_j \varepsilon_j \leq -\sqrt{\varepsilon_j},$$

so its unique slope is $H_j = \varepsilon_j^{-1/2} = 2^j$. Thus nodewise Bellman-attaining slopes diverge although prices and full martingale classes converge. Adjoining $y = \pm 1$ at each bad node restores radius one. For the four successors $\{\pm \varepsilon_j, \pm 1\}$, the optimal slope is

$$H_j^{\text{aug}} = \frac{1 + \sqrt{\varepsilon_j}}{1 + \varepsilon_j} \leq \frac{1 + \sqrt{2}}{2},$$

which makes the repair explicit. □

7.3 Radius-preserving node augmentation

The repair is local and finite. At each retained node one adjoins a bounded number of allowed successors whose increments span a cross-polytope, which restores an interior ball of radius comparable to the continuum radius; the resulting trees remain nested subtrees of Ω and still recover the martingale class.

Lemma 7.4 (Ball-preserving node augmentation). *Let $d_n = \dim D_n$ and suppose $d_n \geq 1$. For every $n < N$ and every $h \in \mathbf{H}_n$ there is a finite set $W_n(h) \subseteq \Gamma_n(h)$ with $\#W_n(h) \leq 2d_n(d_n + 1)$ such that*

$$\{z \in D_n : |z| \leq \kappa_n / \sqrt{d_n}\} \subseteq \text{conv}\{\Delta S_n(h, e) : e \in W_n(h)\}. \quad (72)$$

If $d_n = 0$, so that every increment at date n vanishes and $D_n = \{0\}$, take $W_n(h)$ to be any single successor; the only admissible holding is $H_n = 0$, which contributes nothing to a gain and nothing to a slope bound, so that date may be omitted from both.

Proof. The degenerate case is the parenthetical statement, so assume $d_n \geq 1$ and fix an orthonormal basis $u_1, \dots, u_{d_n}$ of D_n . For each index i and each sign, the point $\pm \kappa_n u_i$ lies in the κ_n -ball of D_n , hence in $\text{conv}\{\Delta S_n(h, e) : e \in \Gamma_n(h)\}$ by (38). Caratheodory's theorem in the d_n -dimensional space D_n writes each of these $2d_n$ points as a convex combination of at most $d_n + 1$ increments; let $W_n(h)$ collect the corresponding successors. Then the displayed convex hull contains every $\pm \kappa_n u_i$, hence contains their convex hull, the scaled cross-polytope $\kappa_n \text{conv}\{\pm u_i\}$. The distance from the origin to each facet of $\text{conv}\{\pm u_i\}$ is $1/\sqrt{d_n}$, so that cross-polytope contains the Euclidean ball of radius $\kappa_n/\sqrt{d_n}$. $\square$

Proposition 7.5 (Augmented cubature trees carry a uniform radius). *Fix one choice of $W_n(h)$ per history. Let Ω_m^{aug} be obtained from Ω_m of Proposition 7.2 by adjoining, at every retained node h of every level $n < N$, the successors $W_n(h)$, and recursing on the adjoined nodes. Then Ω_m^{aug} is a finite nested subtree of Ω , the conclusion (70) continues to hold with $\mathcal{M}_m^{\text{aug}}$ in place of $\mathcal{M}_m$, and every node of Ω_m^{aug} with $d_n \geq 1$ satisfies the relative-interior condition (38) with the radius $\kappa_n/\sqrt{d_n}$, which does not depend on m . At a node with $d_n = 0$, every increment vanishes, the designated set $W_n(h)$ is a singleton, and the only admissible holding is zero.*

Proof. If $d_n \geq 1$, each node acquires at most $2d_n(d_n + 1)$ extra successors; if $d_n = 0$, it acquires at most the single designated successor. Thus the branching stays finite at each of the N levels and the tree is finite; fixing $W_n(h)$ per history keeps the family nested. Every adjoined edge lies in Γ_n at that exact history, so $\Omega_m^{\text{aug}} \subseteq \Omega$ and $\mathcal{M}_m^{\text{aug}} \subseteq \mathcal{M}_\Omega$. Each cubature law $\hat{Q}^{m,k}$ is still a martingale law on the larger tree, assigning the adjoined successors zero mass, so $\mathcal{M}_m^{\text{aug}} \supseteq \mathcal{M}_m \ni \hat{Q}^{m,k}$ and the two-sided estimate in the proof of Proposition 7.2 is unchanged. The radius assertion is (72). $\square$

7.4 Values, margins, and lifted certificates

With law recovery and, on the augmented sequence, a uniform radius in hand, we now pass to the limit in the objects that matter for execution: robust values, separation margins and directions, the uniform bound on tree-optimal holdings, and finally the pointwise lifting of a finite-tree certificate to the continuum.

Theorem 7.6 (Convergence of values and separating portfolios). *Let $(\Omega_m, \mathcal{M}_m)$ denote either the sequence constructed in Proposition 7.2 or its augmented version from Proposition 7.5, with the superscript aug suppressed in the latter case. Then, for every $f \in C(\Omega)$,*

$$\max_{Q \in \mathcal{M}_m} E_Q f \longrightarrow \max_{Q \in \mathcal{M}_\Omega} E_Q f, \quad \min_{Q \in \mathcal{M}_m} E_Q f \longrightarrow \min_{Q \in \mathcal{M}_\Omega} E_Q f. \quad (73)$$

Fix a nonempty, noncrossed finite quote list $J = (a_1, \dots, a_q)$, so $q \geq 1$, and put

$$\begin{aligned} A_m &= \{(E_Q g_{a_i})_{i=1}^q : Q \in \mathcal{M}_m\}, \\ A &= \{(E_Q g_{a_i})_{i=1}^q : Q \in \mathcal{M}_\Omega\}, \\ B &= \prod_{i=1}^q [\underline{q}_{a_i}, \bar{q}_{a_i}], \end{aligned}$$

and define the normalized strict-separation values

$$\gamma_m = \max_{|\theta|=1} \left\{ \min_{z \in A_m} \theta \cdot z - \sup_{y \in B} \theta \cdot y \right\}, \quad (74)$$

$$\gamma = \max_{|\theta|=1} \left\{ \min_{z \in A} \theta \cdot z - \sup_{y \in B} \theta \cdot y \right\}. \quad (75)$$

Then $A_m \rightarrow A$ in Euclidean Hausdorff distance and $\gamma_m \rightarrow \gamma$. If $A \cap B = \emptyset$, then $\gamma > 0$, all sufficiently large finite trees are infeasible, and any sequence of maximizers θ_m in (74) has a subsequence converging to a maximizer in (75). Along every such subsequence the static directions converge to a continuum optimizer, while the robust margins and optimal dynamic completion values converge along the full sequence. No convergence of the hedge ratios is asserted.

Proof. For continuous f , the map $Q \mapsto E_Q f$ is uniformly continuous on the compact space $\mathcal{P}(\Omega)$. Hausdorff convergence in (70) therefore gives both limits in (73).

The vector map

$$T(Q) = (E_Q g_{a_1}, \dots, E_Q g_{a_q})$$

is also uniformly continuous. Hence the Hausdorff distance between $T(\mathcal{M}_m) = A_m$ and $T(\mathcal{M}_\Omega) = A$ tends to zero. For $|\theta| = 1$,

$$\left| \min_{z \in A_m} \theta \cdot z - \min_{z \in A} \theta \cdot z \right| \leq d_H(A_m, A).$$

Thus the objectives inside (74) converge uniformly on the unit sphere, which proves $\gamma_m \rightarrow \gamma$ and the assertion about maximizers. If A and B are disjoint, compact convex strict separation shows $\gamma > 0$. The remaining claims follow from positivity of γ_m for all large m and finite-tree duality. More explicitly, the map $\theta \mapsto \sum_i \theta_i g_{a_i} - c_J(\theta)$ is continuous from the unit sphere to $C(\Omega)$ in the uniform norm. Hence (73) is uniform over these separator payoffs, proving convergence of their dynamic completion values even when $\theta = \theta_m$ varies with m . $\square$

The convergence in Theorem 7.6 is quantitative. Write $\varepsilon_m := d_H^w(\mathcal{M}_m, \mathcal{M}_\Omega)$ for either sequence; by the proof of Proposition 7.2, $\varepsilon_m \leq 2/m$, and the same bound holds for the augmented sequence because $\mathcal{M}_m \subseteq \mathcal{M}_m^{\text{aug}} \subseteq \mathcal{M}_\Omega$. For a finite combination $f_J = \sum_{j \leq J} c_j \varphi_j$ of the test functions in (69) put $\|f_J\|_w := \sum_{j \leq J} 2^j |c_j|$.

Proposition 7.7 (Explicit value and margin rates). *1. For every $f \in C(\Omega)$ and every finite combination f_J ,*

$$\left| \max_{Q \in \mathcal{M}_m} E_Q f - \max_{Q \in \mathcal{M}_\Omega} E_Q f \right| \leq 2\|f - f_J\|_\infty + \|f_J\|_w \varepsilon_m,$$

and likewise for the minima. In particular the values of payoffs in the span of the test functions converge at rate $O(\varepsilon_m) = O(1/m)$.

2. For the quote list $J = (a_1, \dots, a_q)$, if $\|g_{a_i} - f_J^{(i)}\|_\infty \leq \eta_i$ for finite combinations $f_J^{(i)}$, then with $\delta_m := (\sum_{i \leq q} (2\eta_i + \|f_J^{(i)}\|_w \varepsilon_m)^2)^{1/2}$ the Euclidean Hausdorff distance satisfies $d_H(A_m, A) \leq \delta_m$ and the normalized margins satisfy $|\gamma_m - \gamma| \leq \delta_m$.

Proof. (i) By (69), $|E_P \varphi_j - E_Q \varphi_j| \leq 2^j d_w(P, Q)$, so $|E_P f_J - E_Q f_J| \leq \|f_J\|_w d_w(P, Q)$: the map $Q \mapsto E_Q f_J$ is $\|f_J\|_w$ -Lipschitz for d_w , and maxima of such a map over two sets at Hausdorff distance ε_m differ by at most $\|f_J\|_w \varepsilon_m$. Replacing f_J by f changes each maximum by at most $\|f - f_J\|_\infty$. The minima are handled identically.

(ii) For $Q \in \mathcal{M}_\Omega$ pick $P \in \mathcal{M}_m$ with $d_w(P, Q) \leq \varepsilon_m$; by the argument in (i) applied to each coordinate, $|E_Q g_{a_i} - E_P g_{a_i}| \leq 2\eta_i + \|f_J^{(i)}\|_w \varepsilon_m$, so the point of A attached to Q is within δ_m of a point of A_m ; since $A_m \subseteq A$, $d_H(A_m, A) \leq \delta_m$. For each unit θ , $|\min_{A_m} \theta \cdot z - \min_A \theta \cdot z| \leq \delta_m$, and the maximum over θ in (74) and (75) preserves this bound. $\square$

The theorem deliberately separates value convergence from off-grid execution. The following proposition gives a sufficient, checkable lifting condition for the latter and quantifies the cash buffer.

Proposition 7.8 (Shadow-tree lifting with an explicit buffer). *For the selected finite-tree sequence, write*

$$\mathbf{H}_n^m = \{h_n(\omega^m) : \omega^m \in \Omega_m\}$$

for its level- n raw histories. Suppose there are Borel maps

$$r_{m,n} : \mathbf{H}_n \longrightarrow \mathbf{H}_n^m, \quad n = 0, \dots, N,$$

which are prefix-compatible in the following literal sense: if $h_{n+1} = (h_n, e_{n+1})$, then the level- n prefix of $r_{m,n+1}(h_{n+1})$ is $r_{m,n}(h_n)$. In particular, $r_{m,N}$ maps Ω into Ω_m . For the fixed list J , assume

$$\eta_m = \max_{0 \leq n \leq N} \sup_{h \in \mathbf{H}_n} |S_n(h) - S_n(r_{m,n}(h))| \longrightarrow 0, \quad (76)$$

$$\xi_m = \max_{1 \leq i \leq q} \sup_{\omega \in \Omega} |g_{a_i}(\omega) - g_{a_i}(r_{m,N}(\omega))| \longrightarrow 0. \quad (77)$$

For every m , let a finite-tree separating portfolio have $|\theta_m| = 1$, normalized margin γ_m , and predictable holdings $H_n^m : \mathbf{H}_n^m \rightarrow D_n$, $n < N$, satisfying

$$f_{\theta_m}(\omega^m) + (H^m \cdot S)_N(\omega^m) \geq \gamma_m \quad \text{for every } \omega^m \in \Omega_m, \quad (78)$$

and suppose

$$\sup_m \max_{n < N} \max_{h \in \mathbf{H}_n^m} |H_n^m(h)| \leq L < \infty. \quad (79)$$

Then the lifted predictable strategy $\hat{H}_n^m(h) = H_n^m(r_{m,n}(h))$ satisfies, on every continuum path,

$$f_{\theta_m}(\omega) + (\hat{H}^m \cdot S)_N(\omega) \geq \gamma_m - \sqrt{q} \xi_m - 2NL\eta_m, \quad (80)$$

where $f_{\theta_m} = \sum_i \theta_{m,i} g_{a_i} - c_J(\theta_m)$. In particular, if $\gamma_m \rightarrow \gamma > 0$, the finite-tree portfolio and its lift are pointwise certificates for the continuous representative marks for all large m . Replacing any fixing marks by their actual payment ledgers and making the exact strategy adjustment in (45) gives executable actual-ledger certificates with the same margins.

Proof. Evaluate (78) on the shadow path $r_{m,N}(\omega)$. The static-payoff error is at most $\|\theta_m\|_1 \xi_m \leq \sqrt{q} \xi_m$. Prefix compatibility and (79) give

$$d_{m,n}(h) = S_n(h) - S_n(r_{m,n}(h)), \quad |d_{m,n}(h)| \leq \eta_m,$$

and therefore

$$\begin{aligned} & \left| (\hat{H}^m \cdot S)_N(\omega) - (H^m \cdot S)_N(r_{m,N}(\omega)) \right| \\ & \leq \sum_{n=0}^{N-1} |H_n^m(r_{m,n}(h_n))| (|d_{m,n+1}(h_{n+1})| + |d_{m,n}(h_n)|) \\ & \leq 2NL\eta_m. \end{aligned}$$

Subtracting the two errors proves (80). □

Corollary 7.9 (Quantitative lifting on augmented trees). *Fix a nonempty, noncrossed quote list $J = (a_1, \dots, a_q)$, so $q \geq 1$, and use the augmented sequence. For each maximizing direction θ_m in (74), let $\tilde{H}^m$ be an attained lower-Bellman optimizer for $g_{\theta_m} = \sum_i \theta_{m,i} g_{a_i}$, and set $H^m = -\tilde{H}^m$. Then H^m completes $f_{\theta_m} = g_{\theta_m} - c_J(\theta_m)$ in (78), and may be chosen so that*

$$\sup_{m \geq 1} \max_{n < N} \max_{h \in H_n^m} |H_n^m(h)| \leq L_J := \max_{\substack{n < N \\ d_n \geq 1}} \frac{2\sqrt{q} \max_i \|g_{a_i}\|_\infty \sqrt{d_n}}{\kappa_n}, \quad (81)$$

with $L_J = 0$ when every $d_n = 0$. For each a_i , let K^{a_i} be the bounded carry in (45), taking $K^{a_i} = 0$ when the representative is already the payment ledger, put

$$\|K^{a_i}\|_{\mathcal{H},\infty} := \max_{n < N} \sup_{h \in H_n} |K_n^{a_i}(h)|,$$

and define

$$H^{\text{pay},m} = H^m - \sum_i \theta_{m,i} K^{a_i}.$$

Replacing the representative marks by the actual ledgers and using $H^{\text{pay},m}$ preserves both cost and margin. In particular,

$$\sup_{m \geq 1} \max_{n < N} \max_{h \in H_n^m} |H_n^{\text{pay},m}(h)| \leq L_J + \sqrt{q} \max_i \|K^{a_i}\|_{\mathcal{H},\infty}.$$

If $\gamma_m > 0$, this actual-ledger portfolio is an executable finite-tree certificate. If the shadow-map estimates of Proposition 7.8 also hold, first lifting the representative-mark portfolio and then making the continuum carry adjustment gives the pointwise actual-ledger lower bound

$$\gamma_m - \sqrt{q} \xi_m - 2NL_J \eta_m.$$

Whenever this quantity is positive, the adjusted lift is a genuine executable pointwise continuum certificate. In particular, this holds for all sufficiently large m if $A \cap B = \emptyset$.

Proof. The uniform hedge bound in (79) is not an extra assumption on the augmented sequence. The cubature trees of Proposition 7.2 do not by themselves inherit (38): that hypothesis constrains the continuum successor set $\Gamma_n(h)$, whereas cubature retains only the Caratheodory atoms selected to match finitely many integrals, and at an interior node those atoms come from the cubature laws through that exact history. Their increment hull can therefore be arbitrarily small while the continuum hull still contains the κ_n -ball, and the attained local slope then diverges. Passing to Ω_m^{aug} of Proposition 7.5 repairs this at the cost of finitely many extra successors per node. On that sequence, for each maximizing direction θ_m in (74), select an attained lower-Bellman optimizer $\tilde{H}^m$ and put $H^m = -\tilde{H}^m$. Finite-tree subhedging Proposition 3.5 gives (78). Applying (48) with the radius $\kappa_n/\sqrt{d_n}$ gives

$$\sup_{m \geq 1} \max_{n < N} \max_{h \in H_n^m} |H_n^m(h)| \leq \max_{\substack{n < N \\ d_n \geq 1}} \frac{2\sqrt{q} \max_i \|g_{a_i}\|_\infty \sqrt{d_n}}{\kappa_n} =: L_J < \infty,$$

with the convention $L_J := 0$ when no date has $d_n \geq 1$. The maximum is otherwise over the dates with a nontrivial active space, since a date with $d_n = 0$ admits only $H_n = 0$. This bound is uniform in m , but it must not be obtained by bounding the full separator payoff f_{θ_m} , which contains the possibly large bid-ask cash constant. Instead, write

$$g_{\theta_m} := \sum_{i=1}^q \theta_{m,i} g_{a_i}, \quad f_{\theta_m} = g_{\theta_m} - c_J(\theta_m).$$

Both Bellman envelope operators commute with additive constants:

$$\begin{aligned}\max_e \{u(e) - c - H \cdot a(e)\} &= \max_e \{u(e) - H \cdot a(e)\} - c, \\ \min_e \{u(e) - c - H \cdot a(e)\} &= \min_e \{u(e) - H \cdot a(e)\} - c.\end{aligned}$$

Consequently the optimizing slopes for f_{θ_m} are exactly those for g_{θ_m} . Since $|\theta_m| = 1$,

$$\|g_{\theta_m}\|_\infty \leq \sum_{i=1}^q |\theta_{m,i}| \|g_{a_i}\|_\infty \leq \sqrt{q} \max_i \|g_{a_i}\|_\infty,$$

and the continuation-value sup norms are bounded by the same quantity. Applying (48) to those continuation values gives the displayed constant L_J , independently of the quote-box cash term. Substitution in (80) gives the stated buffer. Finally, (45) shows pathwise that replacing each fixing mark by its actual ledger and subtracting $\sum_i \theta_{m,i} K^{a_i}$ from the certificate holding changes neither cost nor margin. The triangle and Cauchy–Schwarz inequalities give the displayed bound for $H^{\text{pay},m}$. What remains genuinely hypothetical is the off-grid part: without a shadow map satisfying (76)–(77), or an equivalent cellwise verification, a grid inequality is still not an off-grid hedge. $\square$

7.5 Main compact-state theorem

The preceding results close the finite-to-continuum loop. The following main theorem states the continuum certificate and the finite-tree approximation in one place; its clauses are proved in the component results cited below.

Theorem 7.10 (Compact finite-calendar robust certificates). *Assume:*

- (C1) *the recursively generated history spaces and continuous compact successor correspondences of (37);*
- (C2) *one common-numeraire, cum-dividend, finite executable price vector, with every fixing-to-payment carry represented as in (45);*
- (C3) *the uniform active-space relative-interior condition (38);*
- (C4) *the saturated rectangular dual class $\mathcal{M}_\Omega$ of all prediction-set martingale laws, rather than only a preferred model library;*
- (C5) *continuous normalized quoted-claim representatives, with every representative-to-actual-ledger conversion supplied by (45), and with each synthetic separating potential reduced to finitely many listed claims as in (34).*

For the tree-law and value conclusions, take the augmented sequence constructed in Lemma 7.1 and Propositions 7.2 and 7.5; write it as $(\Omega_m, \mathcal{M}_m)$, suppressing the superscript aug. This is a proved consequence, not an additional recovery assumption. On that sequence, for each fixed nonempty, noncrossed quote list and each sequence of maximizing directions, bounded Bellman-selected optimal dynamic completions satisfying (79) exist by Lemma 7.4 and Corollary 7.9; this is not a bound on arbitrary optimizing holdings. Only for the optional direct lifting of the finite-tree strategies themselves, impose in addition the shadow-map estimates (76)–(77) of Proposition 7.8; call that remaining off-grid condition (C6). Then all of the following conclusions hold.

- (A) *Every continuous payoff has exact pointwise superhedging and subhedging duality with no gap; the model extrema and bounded Borel-strategy extrema are attained.*

- (B) *The Bellman value admits the pathwise optional decomposition (53); continuous predictable ϵ -hedgies exist, and the continuous-payoff superhedging cone is norm closed with polar $\mathcal{M}_\Omega$.*
- (C) *Empty intersection of an arbitrary continuum bid-ask quote book with $\mathcal{M}_\Omega$ is equivalent to a finite executable static portfolio, contractual carry, and bounded predictable dynamic strategy having strictly negative initial cost and strictly positive payoff on every prediction path. The strategy for the continuous fixing-mark representatives can be continuous; after conversion to actual ledgers it remains bounded Borel and is continuous whenever the carry strategies are continuous.*
- (D) *Along the augmented cubature trees, robust prices converge for every continuous payoff. For each fixed nonempty, noncrossed finite quote list, separation margins and optimal dynamic completion values converge, and every sequence of maximizers in (74) has continuum-optimal cluster points. Full direction convergence is asserted only under uniqueness. If $A \cap B = \emptyset$, then $\gamma > 0$, and all sufficiently large finite-tree maximizers, equipped with the bounded Bellman-selected completions of Corollary 7.9, give executable finite-tree certificates after the exact carry adjustment. Under the optional shadow-tree estimates, lifting the continuous representatives and then applying the continuum carry adjustment gives pointwise actual-ledger certificates for all sufficiently large m , with the explicit error buffer (80). Quantitative value, margin, and augmented-tree lifting bounds are given in Proposition 7.7 and Corollary 7.9.*

Proof. Part (A) and the optional-decomposition assertion in (B) are Theorem 6.6; the closed-cone assertion is Corollary 6.7. For part (C), if a contract has bid $\underline{q}$ strictly above ask $\bar{q}$, buy one unit at ask and sell one at bid. The terminal contract payoffs cancel and the initial receipt is $d = \underline{q} - \bar{q} > 0$; placing $d/2$ in the constant account leaves initial cost $-d/2$ and terminal payoff $d/2$. If all intervals are noncrossed, the claim is Theorem 6.8, together with the carry and static spanning identities. Part (D) is Propositions 7.2, 7.7 and 7.8, Theorem 7.6, and Corollary 7.9. $\square$

7.6 Finite-level execution of expectation-level separators

The preceding convergence results explain when finite-level separators are reliable. A separate point is that each strict finite-level separator can already be executed, without taking a limit of strategies. Let m index a refining family of separators, and expand each finite piecewise-linear optimizer by (34) into actually quoted mass, forward, call, or put legs. Include every deterministic notional, accrual, initial-mass, and state-price normalization factor. Let g_m be the resulting continuous fixing expression and p_m its long-at-ask, short-at-bid acquisition cost.

For the direct inference below, one must establish a positive lower margin over the saturated law class used by the trading theorem. We therefore take

$$\mathcal{M} = \mathcal{Q}_{\mathbb{T}} \neq \emptyset \quad \text{on a finite tree,} \quad \text{or} \quad \mathcal{M} = \mathcal{M}_\Omega \quad \text{on a compact prediction set.}$$

If $\mathcal{Q}_{\mathbb{T}} = \emptyset$, the finite-tree Farkas alternative already supplies a pure cash/dynamic certificate, so no subhedging argument is needed. A positive margin computed over a strict, preferred subclass of models is by itself only model rejection, unless a separate argument transfers the positive lower bound to the saturated class; equality of the two relevant lower values is one sufficient transfer. Suppose that

$$\inf_{Q \in \mathcal{M}} E_Q g_m - p_m \geq \delta_m > 0. \tag{82}$$

Assume that g_m is exactly carry-admissible, with actual payment ledger $\tilde{g}_m$ and predictable carry K_m^{car} ; in the compact case require this carry to be bounded and Borel-predictable:

$$\tilde{g}_m = g_m + (K_m^{\text{car}} \cdot X)_N \quad \text{pathwise.}$$

Choose

$$p_m < z_m < \inf_{Q \in \mathcal{M}} E_Q g_m.$$

Finite-tree or compact-state robust subhedging of the continuous fixing mark, followed by the exact carry, supplies a predictable H^m with

$$z_m + (H^m \cdot X)_N \leq \tilde{g}_m \quad \text{on every admissible path.} \quad (83)$$

Here H^m is the fixing-mark subhedge plus K_m^{car} . Consequently

$$\tilde{g}_m - (H^m \cdot X)_N - p_m \geq z_m - p_m > 0 \quad \text{on every admissible path.} \quad (84)$$

Buying the actual claim for p_m and shorting the subhedge is therefore an executable certificate. Its quoted-leg signs are exactly the long-at-ask/short-at-bid signs recovered by the Farkas alternative.

Proposition 7.11 (Finite-level execution principle). *Suppose (82) holds for all sufficiently large m . If each corresponding finite fixing expression is exactly carry-admissible, with bounded Borel-predictable carry in the compact case, and the robust subhedging equality applies to it over the same saturated law class (or the strict bound has otherwise been validly transferred to that class), then every sufficiently large m supplies an executable robust arbitrage certificate. No convergence of the strategies H^m is needed.*

Proof. For each fixed m , choose z_m strictly between market cost and robust lower value. Robust subhedging and the pathwise carry identity give (83); subtracting that inequality from the purchased claim gives (84). The construction is pointwise and takes no limit of stochastic integrals. $\square$

8 Boundaries and scope

The following failures are structural rather than technical omissions. They also identify which part of the certificate pipeline breaks when a standing hypothesis is removed.

8.1 The law class: saturation and rectangularity

For a restricted family of models $\mathcal{Q}_0$, an inequality $\inf_{Q \in \mathcal{Q}_0} E_Q G > p$ by itself only rejects that family. To infer a market arbitrage one must additionally establish the same strict lower bound over the declared full prediction-set martingale class and invoke a robust subhedging theorem that converts it into a pathwise trading inequality. Weak density of $\mathcal{Q}_0$ in $\mathcal{M}_\Omega$ is one sufficient transfer when G is continuous, but it must be proved rather than inferred from statistical fit. For example, let $\Omega = \{0, 1\}$, let the dynamic asset be constant, and let $g(\omega) = \omega$. The preferred singleton containing only the Bernoulli(1/2) law cannot calibrate the price 1/4, but the Bernoulli(1/4) martingale law on the same prediction set does. Hence this non-dense preferred-class separation is model rejection, not pathwise arbitrage. Working directly with the saturated class in (42) avoids a separate value-identification step.

Rectangularity is a separate requirement. If a law class is not stable under nodewise pasting, backward recursion can combine kernels from different global laws and optimize over laws that were never admissible. Let the two intermediate histories be A, B , let the terminal mark be $z \in \{0, 1\}$, and let the traded asset be constant. For $0 \leq \lambda \leq 1$, define

$$Q_\lambda(A, 0) = (1 - \lambda)/2, \quad Q_\lambda(A, 1) = \lambda/2,$$

$$Q_\lambda(B, 0) = \lambda/2,$$

$$Q_\lambda(B, 1) = (1 - \lambda)/2.$$

The compact convex class $\mathcal{Q} = \{Q_\lambda : 0 \leq \lambda \leq 1\}$ fixes the first-stage probabilities at $1/2, 1/2$ and imposes $Q(z = 1 \mid A) + Q(z = 1 \mid B) = 1$. For $F(A, z) = F(B, z) = z$,

$$\inf_{Q \in \mathcal{Q}} E_Q F = \frac{1}{2}.$$

At either node, however, the projected continuation kernels form the whole Bernoulli family. Nodewise lower recursion selects δ_0 after both histories and returns zero. The pasted law

$$Q^{00} = \frac{1}{2}\delta_{(A,0)} + \frac{1}{2}\delta_{(B,0)}$$

does not belong to $\mathcal{Q}$, because membership would require simultaneously $\lambda = 0$ and $\lambda = 1$. Thus recursion has enlarged the class and changed the value. The compact-state theorem therefore works with the full rectangular kernel class or an explicitly declared rectangular hull; compactness alone does not imply rectangularity.

8.2 Attainment and static execution

The condition $0 \in \text{conv } K$ is sufficient for affine attainment on a finite successor set because the envelope is polyhedral. It is not sufficient on a compact continuum. Take $K = [0, 1]$, current price zero, and $f(y) = -\sqrt{y}$. The only probability on K with mean zero is δ_0 , so the primal lower value is zero. An attaining affine minorant at that value would require a finite h with

$$hy \leq -\sqrt{y} \quad (0 < y \leq 1),$$

which is impossible. Values below zero are feasible with slopes whose magnitude diverges as the value approaches zero. Thus boundary membership can give value duality without an attained bounded hedge; the relative- interior radius rules out precisely this failure. Across arbitrary positive-dimensional, pointwise-interior one-step families, a common balance bound also gives the quantitative converse in the sense of [Theorem 6.4](#).

Static execution is finite-strike. A continuum bounded-Lipschitz potential may admit a formal representation by a distribution of calls, but it is not an executable finite portfolio unless that strike continuum is actually traded with the necessary integrability and tail control. This paper executes finite piecewise-linear potentials through [\(34\)](#); a limiting potential is only a value-limit object.

Contract domains and filtration are equally structural. If an annuity or settlement denominator approaches zero without the required inverse-moment control, the carried payoff may leave the pricing domain even when the potentials themselves are bounded. If a latent regime label is used in a hedge before it is observed by the market filtration, the resulting strategy is clairvoyant. The results use the shared raw canonical filtration and only the contract-compatible domains explicitly declared with the traded claims.

8.3 What convergence does not provide

The law recovery [\(70\)](#) controls integrals of continuous payoffs and hence robust values. It does not control a selected optimizer when affine minorants are nonunique, nor stochastic integrals of oscillating or exploding deltas. The radius-preserving augmentation supplies uniform holdings, and [Proposition 7.8](#) adds the off-grid estimate needed to lift a tree inequality. Alternatively, one reconstructs a continuum hedge from [Theorem 6.6](#); no limit of finite policies is then claimed.

Nor does ordinary path convergence stabilize stochastic gains. For Brownian motion, the left-step approximation converges uniformly in probability while the quadratic variation of the difference does not vanish, so the approximation is not Emery convergence; a small hidden noise can likewise be large in an adapted topology when later observations reveal earlier hidden states. These phenomena are classical (Duffie and Protter, 1992; Kurtz and Protter, 1991; Hubalek and Schachermayer, 1998). The present discretization theorem concerns values, margins, and a direct pointwise lifting buffer, not Emery convergence of strategies.

Finally, compactness of a family of laws does not imply compactness of the union of their supports. A noncompact application therefore needs a declared compact prediction core, a proved tail-localization argument, or an analytic-set/quasi-sure theorem with its own admissibility and closedness proof.

8.4 Finite dimension, continuous time, and open questions

For compact successor sets, a uniform norm ball in the closed convex hull of increments is incompatible with a genuinely infinite-dimensional active Banach space. The theorem allows a continuum of states and statically quoted contracts, but only a finite dynamically traded menu on a fixed calendar. Growing the maturity menu or trading a cylindrical curve requires a different compactness mechanism.

In particular, the compact-state theorem does not assert continuous-time self-financing, predictable uniform tightness, convergence in the Emery topology, closedness of stochastic integrals under refinement, or an optional decomposition for cylindrical curve trading. A continuous-time extension would need:

1. predictable-uniform-tightness or an equivalent characteristic compactness condition for the approximating martingales;
2. closedness of terminal gains under an Emery- or semimartingale-stable topology, including post-fixing carry;
3. a robust optional decomposition on the nondominated prediction set of the extended menu;
4. convergence, or at least compactness and valid cluster points, of the finite-calendar hedges; and
5. preservation of the contractual settlement identities in the continuous-time limit.

Single-reference-law stability results include (Kardaras, 2024; De Donno and Pratelli, 2005; Alberti and Karbach, 2026); related stochastic-dimension arbitrage and optional-decomposition foundations include Bayraktar et al. (2024). The nondominated growing-menu problem is open.

For a fixed finite panel, a relevant benchmark is the nondominated optional decomposition of Nutz (2015, Theorem 2.4), in the continuous-time superhedging line beginning with volatility uncertainty for measurable claims (Neufeld and Nutz, 2013). Its dominating-diffusion hypothesis requires care with multidimensional jumps. Rodrigues identifies a gap in the proof's identification step, not a counterexample to the theorem: the coordinatewise condition does not suffice when the covariance density $c^P = dC^P / d \text{Tr}(C^P)$ is singular (Rodrigues, 2025, Remark 3.3, Example 3.9). The strengthened condition

$$(|x|^2 \wedge 1) * \nu^P \ll \mathbf{1}_{\{\det(c^P) > 0\}} \text{Tr}(C^P)$$

is introduced in Rodrigues (2025, Lemma 3.1, Definition 3.2), and the optional decomposition is

reproved in [Rodrigues \(2025, Theorem 3.7\)](#). A growing-menu extension of that fixed-panel theory would additionally be required.

Two questions already arise in finite time. First, the active subspaces D_n are indexed by date only; history-dependent spaces $D_n(h)$ would require regularity of the varying subspaces as well as a cellwise uniform relative-interior condition. Second, the shadow-map estimates of [Proposition 7.8](#) are sufficient for off-grid execution but are not shown necessary; the paper does not characterize all prediction sets on which finite-tree hedges lift.

9 Conclusion

In the route developed here, an expectation-level quote inconsistency becomes an executable certificate only after three interfaces are secured: the separator has a positive robust margin over the full saturated prediction-set martingale class, established directly or transferred by a separate valid argument, its static terms are finite traded portfolios carried pathwise to their contractual ledgers, and its robust lower value is attained by a bounded predictable hedge. Under assumptions (C1)–(C5) of [Theorem 7.10](#), the uniform local interior radius supplies the last interface. Across arbitrary indexed families of positive-dimensional, pointwise-interior one-step geometries, a common balance bound also gives the quantitative converse from uniformly bounded least-norm optimizers; without a common balance bound the converse fails. When calibration fails, the result is a finite executable certificate with an attained bounded dynamic completion, valid on every declared path, even when the original quote book is uncountable.

Computation adds a further requirement. Martingale cubature can recover the full law class and every continuous-payoff value while destroying the local radius and sending optimal hedge ratios to infinity. Finite node augmentation preserves value convergence for every continuous payoff and margin convergence for each fixed nonempty, noncrossed finite quote list, while restoring a uniform radius and bounded Bellman selections. If an infeasible quote book contains a crossed quote, a cash-funded static round trip is immediate. Otherwise the book first supplies a noncrossed finite witness sub-book, whose finite-tree margins are positive for all sufficiently large trees. Under the stated shadow-map estimates those certificates lift whenever the explicit buffered margin remains positive. Faithful discretization must therefore approximate the geometry of hedging, not only the prices generated by the laws. Continuous-time and growing-menu extensions require additional closedness and stochastic-integral stability and remain outside the present theorem.

A Contract-level instances of exact carry

This appendix records representative interest-rate contracts covered by [Lemma 4.1](#). All dates belong to the marked master calendar, all contract amounts are finite at the determination node, and every carried price process is one of the stopped traded components declared in [Section 4](#).

For a caplet or floorlet determined at θ , let C_θ be its $\mathcal{F}_\theta$ -measurable contractual currency amount paid at p under convention c . The fixing mark and payment ledger are

$$f = Z_\theta^c(p)C_\theta, \quad g = \zeta_p^c C_\theta. \quad (85)$$

After θ , hold C_θ units of the c -collateral unit payment claim $M^{c,p}$. The diagonal identity $M_{n(p)}^{c,p} = \zeta_p^c$ makes this exactly [Lemma 4.1](#). In particular, a payer caplet with deterministic notional n , accrual δ , strike K , and published fixing L_θ has $C_\theta = n\delta(L_\theta - K)^+$; a floorlet replaces the last factor by $(K - L_\theta)^+$.

For a cash-settled payer swaption, let e be exercise, $\sigma \geq e$ the payment date, $\widehat{S}_e$ the contractual settlement rate, and $H_e^a > 0$ the convention- a , strike-independent contractual cash multiplier. Assume $\widehat{S}_e$ and H_e^a are finite and $\mathcal{F}_e$ -measurable. Its payment amount is $C_e = nH_e^a(\widehat{S}_e - K)^+$, and the preceding display applies with $\theta = e$ and $p = \sigma$. The payment date, not the rate determination date, selects the carried state-price claim.

For a matched-collateral physically settled payer swaption, let $\overline{F}_e$ be the common-unit exercise mark of a finite traded floating-leg portfolio and let $\overline{A}_e > 0$ be the corresponding mark of the finite traded portfolio delivering one unit of the fixed coupon schedule, and let $\overline{A}_0 > 0$ be that portfolio's deterministic initial mark. Both exercise marks are finite and $\mathcal{F}_e$ -measurable, their constituent payment claims are stopped components of S , and $S_e := \overline{F}_e/\overline{A}_e$ is the exercise-date par rate. Write the exercise mark as

$$f = (\overline{F}_e - K\overline{A}_e)^+.$$

With $I_e = \mathbf{1}_{\{\overline{F}_e > K\overline{A}_e\}}$, positive homogeneity gives

$$f = I_e(\overline{F}_e - K\overline{A}_e).$$

Both legs on the right are finite signed linear combinations of the primitive coupon-price and state-price claims. On exercise, hold I_e times the floating-leg portfolio and $-I_e K$ times the fixed-leg portfolio through their respective payment dates. The sum of their stopped payment values is the actual delivered-swap ledger g , and (32) follows term by term.

A declared cross-CSA linear liquidation requires a separate statement. Let o be the option collateral convention and u the delivered-swap convention. Assume $\zeta_e^o, \zeta_e^u > 0$, and let $\overline{F}_e^u$ and $\overline{A}_e^u > 0$ be the common-unit marks of the finite delivered floating and fixed-leg portfolios under u . Define $S_e^u = \overline{F}_e^u/\overline{A}_e^u$, the currency annuity $A_e^u = \overline{A}_e^u/\zeta_e^u$, and

$$W_e^{o,u} = \zeta_e^o A_e^u = \frac{\zeta_e^o}{\zeta_e^u} \overline{A}_e^u.$$

An agreement which cash-liquidates the delivered swap at exercise has the actual exercise-date common-unit payoff

$$g = W_e^{o,u}(S_e^u - K)^+.$$

It is already the payment ledger, so its post-exercise carry is zero. If the switch instead leaves future delivery legs, the entire linear switch claim—not merely the two separate factors ζ^o/ζ^u and $\overline{A}^u$ —must have a declared finite-leg traded-price representation to which Lemma 4.1 applies. Separate tradability of two factors does not replicate their product. Optional collateral choice and nonlinear funding remain outside the theorem.

AI-use disclosure

The author used Anthropic Claude Code and OpenAI Codex as interactive research and writing tools. They assisted with exploratory discussion, testing and refinement of ideas, literature and source organization, code development and verification, mathematical error checking, and editorial revision.

References

Beatrice Acciaio, Mathias Beiglböck, Friedrich Penkner, and Walter Schachermayer. A model-free version of the fundamental theorem of asset pricing and the super-replication theorem. *Mathematical Finance*, 26(2):233–251, 2016. doi: 10.1111/mafi.12060.

- Riccardo Alberti and Sven Karbach. Hedging maturity-specific risk in forward curve derivatives under stochastic volatility, 2026. arXiv:2606.28891.
- Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, third edition, 2006.
- Daniel Bartl, Michael Kupper, and Ariel Neufeld. Pathwise superhedging on prediction sets. *Finance and Stochastics*, 24(1):215–248, 2020. doi: 10.1007/s00780-019-00412-4.
- Christian Bayer and Josef Teichmann. The proof of Tchakaloff's theorem. *Proceedings of the American Mathematical Society*, 134(10):3035–3040, 2006. doi: 10.1090/S0002-9939-06-08249-9.
- Erhan Bayraktar and Matteo Burzoni. On the quasi-sure superhedging duality with frictions. *Finance and Stochastics*, 24(1):249–275, 2020. doi: 10.1007/s00780-019-00411-5.
- Erhan Bayraktar and Yuchong Zhang. Fundamental theorem of asset pricing under transaction costs and model uncertainty. *Mathematics of Operations Research*, 41(3):1039–1054, 2016. doi: 10.1287/moor.2015.0767.
- Erhan Bayraktar, Donghan Kim, and Abhishek Tilva. Arbitrage theory in a market of stochastic dimension. *Mathematical Finance*, 34(3):847–895, 2024.
- Mathias Beiglböck and Marcel Nutz. Martingale inequalities and deterministic counterparts. *Electronic Journal of Probability*, 19(95):1–15, 2014. doi: 10.1214/EJP.v19-3270.
- Mathias Beiglböck, Pierre Henry-Labordère, and Friedrich Penkner. Model-independent bounds for option prices—a mass transport approach. *Finance and Stochastics*, 17(3):477–501, 2013. doi: 10.1007/s00780-013-0205-8.
- Mathias Beiglböck, Marcel Nutz, and Nizar Touzi. Complete duality for martingale optimal transport on the line. *The Annals of Probability*, 45(5):3038–3074, 2017. doi: 10.1214/16-AOP1131.
- Bruno Bouchard and Marcel Nutz. Arbitrage and duality in nondominated discrete-time models. *The Annals of Applied Probability*, 25(2):823–859, 2015. doi: 10.1214/14-AAP1011.
- Bruno Bouchard and Marcel Nutz. Consistent price systems under model uncertainty. *Finance and Stochastics*, 20(1):83–98, 2016. doi: 10.1007/s00780-015-0286-7.
- Bruno Bouchard, Shuoqing Deng, and Xiaolu Tan. Superreplication with proportional transaction cost under model uncertainty. *Mathematical Finance*, 29(3):837–860, 2019. doi: 10.1111/mafi.12197.
- Matteo Burzoni. Arbitrage and hedging in model-independent markets with frictions. *SIAM Journal on Financial Mathematics*, 7(1):812–844, 2016. doi: 10.1137/15M1053013.
- Matteo Burzoni, Marco Frittelli, and Marco Maggis. Universal arbitrage aggregator in discrete-time markets under uncertainty. *Finance and Stochastics*, 20(1):1–50, 2016. doi: 10.1007/s00780-015-0283-x.
- Matteo Burzoni, Marco Frittelli, and Marco Maggis. Model-free superhedging duality. *The Annals of Applied Probability*, 27(3):1452–1477, 2017. doi: 10.1214/16-AAP1235.

- Matteo Burzoni, Marco Frittelli, Zhaoxu Hou, Marco Maggis, and Jan Obłój. Pointwise arbitrage pricing theory in discrete time. *Mathematics of Operations Research*, 44(3):1034–1057, 2019. doi: 10.1287/moor.2018.0956.
- Laurence Carassus, Jan Obłój, and Johannes Wiesel. The robust superreplication problem: A dynamic approach. *SIAM Journal on Financial Mathematics*, 10(4):907–941, 2019. doi: 10.1137/18M1235934.
- Patrick Cheridito, Michael Kupper, and Ludovic Tangpi. Duality formulas for robust pricing and hedging in discrete time. *SIAM Journal on Financial Mathematics*, 8(1):738–765, 2017. doi: 10.1137/16M1064088.
- Samuel N. Cohen, Christoph Reisinger, and Sheng Wang. Detecting and repairing arbitrage in traded option prices. *Applied Mathematical Finance*, 27(5):345–373, 2020. doi: 10.1080/1350486X.2020.1846573.
- Laurent Cousot. Conditions on option prices for absence of arbitrage and exact calibration. *Journal of Banking and Finance*, 31(11):3377–3397, 2007. doi: 10.1016/j.jbankfin.2007.04.006.
- Mark H. A. Davis and David G. Hobson. The range of traded option prices. *Mathematical Finance*, 17(1):1–14, 2007. doi: 10.1111/j.1467-9965.2007.00291.x.
- Marzia De Donno and Maurizio Pratelli. A theory of stochastic integration for bond markets. *The Annals of Applied Probability*, 15(4):2773–2791, 2005. doi: 10.1214/105051605000000548.
- Yan Dolinsky and H. Mete Soner. Duality and convergence for binomial markets with friction. *Finance and Stochastics*, 17(3):447–475, 2013. doi: 10.1007/s00780-012-0192-1.
- Yan Dolinsky and H. Mete Soner. Martingale optimal transport and robust hedging in continuous time. *Probability Theory and Related Fields*, 160(1–2):391–427, 2014a. doi: 10.1007/s00440-013-0531-y.
- Yan Dolinsky and H. Mete Soner. Robust hedging with proportional transaction costs. *Finance and Stochastics*, 18(2):327–347, 2014b. doi: 10.1007/s00780-014-0227-x.
- Darrell Duffie and Philip Protter. From discrete- to continuous-time finance: Weak convergence of the financial gain process. *Mathematical Finance*, 2(1):1–15, 1992. doi: 10.1111/j.1467-9965.1992.tb00022.x.
- Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003.
- Hans Föllmer and Dmitry O. Kramkov. Optional decompositions under constraints. *Probability Theory and Related Fields*, 109(1):1–25, 1997. doi: 10.1007/s004400050122.
- Hans Föllmer and Alexander Schied. *Stochastic Finance: An Introduction in Discrete Time*. De Gruyter, Berlin, fifth edition, 2025. doi: 10.1515/97831111045283.
- Stefan Gerhold and Ismail Cetin Gülüm. Consistency of option prices under bid–ask spreads. *Mathematical Finance*, 30(2):377–402, 2020. doi: 10.1111/mafi.12230.
- Gaoyue Guo and Jan Obłój. Computational methods for martingale optimal transport problems. *The Annals of Applied Probability*, 29(6):3311–3347, 2019. doi: 10.1214/19-AAP1481.

- Zhaoxu Hou and Jan Obłój. Robust pricing–hedging dualities in continuous time. *Finance and Stochastics*, 22(3):511–567, 2018. doi: 10.1007/s00780-018-0363-9.
- Friedrich Hubalek and Walter Schachermayer. When does convergence of asset price processes imply convergence of option prices? *Mathematical Finance*, 8(4):385–403, 1998. doi: 10.1111/1467-9965.00060.
- Garud N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.
- Elyès Jouini and Hédi Kallal. Martingales and arbitrage in securities markets with transaction costs. *Journal of Economic Theory*, 66(1):178–197, 1995. doi: 10.1006/jeth.1995.1037.
- Yuri Kabanov and Mher Safarian. *Markets with Transaction Costs: Mathematical Theory*. Springer, Berlin, 2009.
- Olav Kallenberg. *Foundations of Modern Probability*. Springer, third edition, 2021. doi: 10.1007/978-3-030-61871-1.
- Constantinos Kardaras. Stochastic integration with respect to arbitrary collections of continuous semimartingales and applications to mathematical finance. *The Annals of Applied Probability*, 34(3):2566–2599, 2024. doi: 10.1214/23-AAP1942.
- Dmitry O. Kramkov. Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets. *Probability Theory and Related Fields*, 105(4):459–479, 1996. doi: 10.1007/BF01191909.
- Ondřej Kurka. Non-universal families of separable banach spaces. *Studia Mathematica*, 233(2):153–168, 2016. doi: 10.4064/sm8380-4-2016.
- Thomas G. Kurtz and Philip Protter. Weak limit theorems for stochastic integrals and stochastic differential equations. *Annals of Probability*, 19(3):1035–1070, 1991. doi: 10.1214/aop/1176990334.
- Terry Lyons and Nicolas Victoir. Cubature on Wiener space. *Proceedings of the Royal Society of London. Series A*, 460(2041):169–198, 2004. doi: 10.1098/rspa.2003.1239.
- Per Aslak Mykland. Conservative delta hedging. *The Annals of Applied Probability*, 10(2):664–683, 2000. doi: 10.1214/aoap/1019487360.
- Per Aslak Mykland. Financial options and statistical prediction intervals. *The Annals of Statistics*, 31(5):1413–1438, 2003. doi: 10.1214/aos/1065705113.
- Ariel Neufeld and Marcel Nutz. Superreplication under volatility uncertainty for measurable claims. *Electronic Journal of Probability*, 18(48):1–14, 2013. doi: 10.1214/EJP.v18-2358.
- Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
- Marcel Nutz. Superreplication under model uncertainty in discrete time. *Finance and Stochastics*, 18(4):791–803, 2014. doi: 10.1007/s00780-014-0238-7.
- Marcel Nutz. Robust superhedging with jumps and diffusion. *Stochastic Processes and their Applications*, 125(12):4543–4555, 2015. doi: 10.1016/j.spa.2015.07.008.

- Jan Obłój and Johannes Wiesel. A unified framework for robust modelling of financial markets in discrete time. *Finance and Stochastics*, 25(3):427–468, 2021. doi: 10.1007/s00780-021-00454-7.
- Teemu Pennanen. Convex duality in stochastic optimization and mathematical finance. *Mathematics of Operations Research*, 36(2):340–362, 2011. doi: 10.1287/moor.1110.0485.
- Frank Riedel. Financial economics without probabilistic prior assumptions. *Decisions in Economics and Finance*, 38(1):75–91, 2015. doi: 10.1007/s10203-014-0159-0.
- R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- Marco Rodrigues. Robust hedging of american options via aggregated snell envelopes, 2025.
- Alet Roux and Tomasz Zastawniak. American options under proportional transaction costs: Pricing, hedging and stopping algorithms for long and short positions. *Acta Applicandae Mathematicae*, 106(2):199–228, 2009. doi: 10.1007/s10440-008-9290-7.
- Walter Schachermayer. The fundamental theorem of asset pricing under proportional transaction costs in finite discrete time. *Mathematical Finance*, 14(1):19–48, 2004. doi: 10.1111/j.0960-1627.2004.00180.x.
- Vladimir Tchakaloff. Formules de cubatures mécaniques à coefficients non négatifs. *Bulletin des Sciences Mathématiques (2)*, 81:123–134, 1957.