

Non-Convex Market Makers

Coherence without convexity, friction, and parallel composition

Peter Cotton

Working draft v0.1 · August 29, 2026

Abstract

Cost-function market makers are always assumed convex. We show convexity is not what makes them coherent. For a path-independent maker, the exact no-arbitrage condition is a chord condition, that every chord slope of the cost lie in the convex hull of payoffs; convexity enters the standard axiomatics through a separate requirement, monotone information incorporation, and dropping it leaves coherent non-convex makers. Rational flow against such a maker sees the convex envelope: optimal fills land on the contact set, and the non-convexity gap at the starting state passes through to the next trader, so non-convexity costs expressiveness rather than soundness. Bounded incoherence is priced: a proportional fee of at least the arbitrage depth restores no-arbitrage, and by conjugation that fee is exactly a bid-ask spread, with the exact no-trade interval given by fee-widened one-sided chord bounds for any cost whatever. Combining fee-bearing makers is an infimal convolution solved by a single monotone clearing-price root-find; the fees act as an L1 penalty, so participation is sparse and the aggregate supply curve is a consolidated limit order book. Adding a deep quadratic co-quoter is Moreau smoothing, which preserves coherence at every depth while attenuating but never closing the envelope gap. Finally, separation makes model inconsistency financially discoverable: any quoted configuration outside the coherent hull admits a portfolio priced below its worst payoff, so arbitrageurs act as decentralized separation oracles, and coherence is relative to friction and to the cost of finding certificates.

1 Many predictors are already markets

A cost-function market maker posts a potential C over its inventory q of outstanding shares and charges $C(q + s) - C(q)$ for a fill of s shares; prices are gradients of C , and when C is convex the standard theory applies ([Hanson 2003](#); [Abernethy et al. 2013](#)). A run of published identities makes many learning procedures markets in this literal sense. Trading against such a maker is follow-the-regularized-leader with the cost as conjugate regularizer ([Chen and Vaughan 2010](#); [Abernethy et al. 2013](#)), with trader wealth as the learning rate ([Frongillo et al. 2012](#)). A market of Kelly bettors performs Bayesian model averaging with wealths as posterior weights ([Beygelzimer et al. 2012](#)), and equilibrium prices of utility-maximizing agents implement mixtures and products of experts ([Pennock and Wellman 1997](#); [Storkey 2011](#); [Storkey et al. 2012](#); [Barbu and Lay 2012](#)).

Two further identities are used throughout. Precision-weighted estimation of a common mean is a market of data points: each observation a quadratic maker quoting its value with capital equal to its precision (inverse variance), the estimate their merge, since merging makers is the infimal convolution of their costs and liquidity, the inverse of price impact, adds ([Bhaskara](#)

et al. 2023; Barrieu and El Karoui 2005). General least squares with regressors is not this one-quantity merge: observation i contributes the rank-one potential $(y_i - x_i^\top \beta)^2 / (2\sigma_i^2)$ and the potentials add in parameter space, which asks for vector securities or for eliminating internal variables from a network of makers, treated in the companion paper. And a proximal step is a trade against a fee-bearing maker: the proximal operator of $f|\cdot|$ is the soft-threshold, which is the optimal response to a proportional fee (Lemma 4 below), and the prox of any convex g is the response to a unit quadratic maker charging $\frac{1}{2}s^2 + g(s)$, since $\text{prox}_g(x) = \arg \max_s \{xs - \frac{1}{2}s^2 - g(s)\}$; the quadratic term is the maker’s own curvature, and without it the trader’s problem is a bare conjugate, unbounded for $|\mu| > f$ in the ℓ_1 case. Lemma 4 carries the base cost g_q for exactly this reason. The two canonical penalties are then the two market primitives: ridge is a zero-quoting participant with capital λ , lasso is a fee of λ , and the theorem that regularization is robustness to data perturbation (El Ghaoui and Le Bret 1997; Xu et al. 2009) acquires a market reading in which the adversary’s budget is priced rather than assumed.

We ask the converse: which predictors are markets, and what does the market form demand? Sections 2–4 treat one maker: coherence, the biconjugate, and friction. Sections 5–7 treat makers in parallel: fees as spreads, routing, the order book. Section 8 turns cross-market inconsistency into arbitrage, and section 9 states the characterization and its limits. Serial composition, where markets are chained along a factorization, is developed in a companion paper; nothing below depends on it.

What is and is not new here should be said at the outset, because the nearest neighbours are close. The identities of the preceding two paragraphs are used and cited, not claimed. That competing makers operate in parallel, that their aggregate is an infimal convolution, that conjugates add and liquidity adds, and that fees can be levied on the aggregate, are established in Bhaskara et al. (2023); the risk-sharing form of the same law is Barrieu and El Karoui (2005), and cost-function markets, their duality and their online-learning reading are Abernethy et al. (2013) and Chen and Vaughan (2010). Against that background the residue claimed here is: coherence without convexity and the contact-set consequences (§§2–3); the exact scalar routing formula with fees inside the costs, its no-trade bands and sparse participation, and the consolidated-order-book reading (§§5–6); arbitrage depth, the friction characterization, and accessible coherence (§§4, 9); and the synthesis with regularization and incentive closure (§§1, 9). Results that are not new are labelled *Remark* and cited to their sources rather than numbered as contributions: Remark 7 is the aggregation law, used here and not claimed.

Two payoff regimes appear and must not be conflated. In the *bounded* regime a scalar security settles at $\varphi(\omega) = \omega \in [-1, 1]$, vector statements substituting the convex hull of payoff vectors; this is where Proposition 1 bites, where bounded worst-case loss is meaningful, and where sections 2–5 live. In the *unrestricted* regime the security settles at a real-valued quantity with payoff hull $\mathbb{R}$, so the chord condition is vacuous for finite positions and a quadratic maker $C(q) = mq + q^2/(2\lambda)$, whose chord slopes are unbounded, is admissible; this is the estimation regime of the Gaussian fusion of §6 and of Proposition 10. Coherence statements transfer between the regimes only through the hull that defines them. Reference implementations and numerical theorem tests accompany the paper in the `mechanisms` repository.

2 Coherence is a chord condition

Throughout, *arbitrage* means a fill whose profit is strictly positive for every outcome, and *no-arbitrage* the absence of one; this is the strong or sure-loss convention, and it counts a quote on

the boundary of the payoff hull as coherent even though the fill it permits is weakly profitable, breaking even on the outcomes that attain the boundary payoff and gaining on the rest. Readers importing the nonnegative-payoff convention should add relative interiors throughout.

Proposition 1 (no-arbitrage without convexity). *For a path-independent maker (one whose charge depends only on the inventory endpoints) with cost C , no convexity assumed, there is no outcome-independent strictly profitable fill from any state if and only if every chord slope $[C(q + s) - C(q)]/s$ lies in the convex hull of payoffs: for all q, s ,*

$$C(q + s) - C(q) \geq \min_{\omega} \varphi(\omega) s \quad (\text{scalar: } C \text{ is 1-Lipschitz}).$$

Proof. A fill s from state q has sure profit $\min_{\omega} \varphi(\omega) s - [C(q + s) - C(q)]$; positivity for some pair (q, s) is exactly the failure of the displayed inequality, and the inequality for all pairs says every chord slope is supported by the hull. ■

Round trips are refunds for any C , convex or not, because the charge telescopes over closed paths. When the chord condition fails the exploit is the accumulation of sure-profit net positions at states where quotes exit the hull.

Convexity enters the standard axiomatics through *information incorporation*, the marginal-cost monotonicity condition of Abernethy et al. (2013), not through no-arbitrage. The habit of assuming convexity has bundled two different guarantees: monotone price response to flow, and absence of sure-loss opportunities. The axiom sets impose the first, convexity follows, and no-arbitrage is then derived with convexity in hand, so the possibility of a coherent maker with non-monotone quotes never arises. Proposition 1 separates the guarantees: dropping information incorporation while keeping the chord condition leaves a coherent non-convex maker, exhibited numerically in the companion repository. Arbitrage theory without convexity is developed in a different formalism by Lépinette and Tran (2017).

3 The market trades the biconjugate

Assume for this section that C is finite and admits an affine minorant, equivalently that C^* is proper, and let $\hat{C} = C^{**}$ denote the lower convex envelope (Rockafellar 1970) and $g = C - \hat{C} \geq 0$ the gap. Chord coherence alone does not give this. The cost $C(q) = -\alpha|q|$ with $\alpha < 1$ is α -Lipschitz and so chord-coherent for the hull $[-1, 1]$, yet an affine minorant would need slope at most $-\alpha$ and at least α at once; hence $C^* \equiv +\infty$, $C^{**} \equiv -\infty$, the gap is undefined, and at $q = 0$, $\mu = 0$ the trader's supremum $\sup_s \alpha|s|$ is already infinite. Coherence bounds what a single fill can extract for sure; it does not bound speculative profit against a belief.

Proposition 2 (contact and pass-through). *Let C be as above and let a myopic risk-neutral trader have believed mean μ in the hull with $C^*(\mu) < \infty$, facing the maker at state q . Then the trader: (i) has every attained optimal fill landing on the contact set $\{C = \hat{C}\}$, an optimum existing whenever $x \mapsto C(x) - \mu x$ attains its infimum (for instance when it is coercive); and (ii) earns maximal expected profit*

$$\Pi_C(q, \mu) = \Pi_{\hat{C}}(q, \mu) + g(q)$$

as an identity of extended-real suprema, the envelope profit plus the gap at the starting state.

Proof. $\sup_s \mu s - [C(q + s) - C(q)] = \sup_x [\mu x - C(x)] - \mu q + C(q)$, so the profit is $C^*(\mu)$ up to terms in q alone. Since conjugation is invariant under biconjugation, $\hat{C}^* = C^{***} =$

C^* , the two suprema agree whether or not either is attained, which gives (ii) after adding and subtracting $\hat{C}(q)$. For (i), a maximizer of an affine function minus C is a point where an affine minorant touches C , hence touches $\hat{C}$, so any attained optimum lies in the contact set. ■

Attainment is a genuine hypothesis, not a formality. The chord-coherent cost $C(x) = \arctan x$ has lower convex envelope the constant $-\pi/2$ and hence empty contact set; at $\mu = 0$ the supremum $\pi/2$ is approached only as $x \rightarrow -\infty$. Where the relevant contact sits at infinity, the language below about flow landing on contact points and off-contact states being transient describes a limit that no fill realizes.

The maker is value-equivalent to its convex envelope for an unconstrained one-shot risk-neutral optimizer, and no more than that: the optimal endpoints of C are exactly the envelope-optimal endpoints lying in the contact set, so C deletes the envelope's off-contact optima, and partial trades, tie-breaking, capital constraints and noisy flow distinguish the two immediately. Concave stretches are unquotable intermediate *inventory* states that rational flow jumps across. Whoever lands inside one (noise) overpays the gap at the landing state, and by (ii) the next rational trader recoups it, so the maker is a conduit keeping envelope differences over any rational-to-rational span, and off-contact states are transient. The trade-set analogue for constant-function market makers (CFMMs, the decentralized-exchange design) is the canonical concave trading function: an arbitrary invariant is behaviorally equivalent to a concave one (Angeris et al. 2024), with the limits of concavification mapped by Frongillo et al. (2024).

What the gap does to the book is the opposite of a gap in prices. Between contact points $q_1 < q_2$ the envelope is affine with a single supporting slope μ_0 , so the marginal price is constant across the excluded inventory range: as belief crosses μ_0 the optimal inventory jumps from q_1 to q_2 while the price does not move. In the convexified supply correspondence this is an inventory jump of size $q_2 - q_1$ at the supporting price μ_0 , a depth spike rather than a missing price range. Against the original non-convex cost the jump is generally indivisible: for the double well of §6 the full trade from $-c$ to c costs nothing while the half trade from $-c$ to 0 costs αc , so the prefixes of the block are not executable at μ_0 and this is a lumpy endpoint trade, not a conventional divisible level. Non-convexity of C also does not by itself make an implied density bimodal: C^* is convex, so any factor read as $\exp(-C^*)$ is log-concave. Earlier drafts of this paper described the gap as a hole in the book and as a multimodal quote; both readings were wrong, and no claim about bimodal implied densities (Melick and Thomas 1997; Clark and Amen 2017) is made here.

4 Frictions price the remaining failure

The failure mode surviving §2 is chord slopes exiting the hull.

Proposition 3 (fees buy bounded incoherence). *If chord slopes exit the hull by at most ε per unit, i.e. $C(q + s) - C(q) \geq \min_{\omega} \varphi(\omega)s - \varepsilon|s|$, then a proportional fee $f \geq \varepsilon$ restores no-arbitrage.*

Proof. Sure profit with the fee is $\min_{\omega} \varphi(\omega)s - \Delta C - f|s| \leq (\varepsilon - f)|s| \leq 0$. ■

This is the mechanism-level counterpart of a theorem of mathematical finance: under regularity conditions such as sticky paths or full support, arbitrarily small proportional transaction costs restore no-arbitrage for a large class of processes that are arbitrageable frictionlessly, with a consistent price system inside the spread as certificate (Guasoni 2006; Guasoni et al. 2010). In prediction markets the precedent is fees sized to expected arbitrage profit restoring bounded loss under privacy noise (Cummings et al. 2016; Frongillo and Waggoner 2018). Call the smallest ε for which the hypothesis of Proposition 3 holds the cost's *arbitrage depth*: the worst per-unit

excursion of its chord slopes beyond the payoff hull, and so the minimum viable half-spread (the quoted spread being $2f$, since bid and ask sit at $m \mp f$). A venue quoting a wide spread to cover a deeply incoherent cost is arbitrage-free but uninformative in proportion: the market prices the model's incoherence as uncertainty.

5 A linear fee is a bid-ask spread

Why a fee at all? A cost-function maker's charge telescopes, so round trips are free: the maker earns nothing on uninformed flow and cannot recover from volume its adverse-selection losses, the losses to better-informed traders. That no state-dependent cost can charge a round trip, and that a path-dependent volume charge repairs it, is due to Othman and Sandholm (2012); the linear charge $f|s|$ is their volume levy in its simplest convex form, chosen because it conjugates in closed form. Write $g_q(s) = C(q+s) - C(q)$ for a convex C and $m = C'(q)$ for the marginal price.

Lemma 4 (fee–spread duality). *Let $T_f(x) = \text{sign}(x) \max(|x| - f, 0)$ denote the soft-threshold. Then the fee-bearing cost $\tilde{C}_q = g_q + f|\cdot|$ has conjugate*

$$\tilde{C}_q^*(p) = g_q^*(m + T_f(p - m)),$$

zero if and only if $|p - m| \leq f$.

Proof. The conjugate of a sum of closed proper convex functions with overlapping relative interiors of domains is the infimal convolution of the conjugates (Rockafellar 1970, Thm. 16.4), and the conjugate of $f|\cdot|$ is the indicator of $[-f, f]$. Hence $\tilde{C}_q^*(p) = \min_{|u| \leq f} g_q^*(p - u)$, the minimum of the convex function g_q^* over $[p - f, p + f]$. Since $g_q^* \geq 0$ with equality exactly at m , the minimum is attained at the projection of m onto the interval, which is $m + T_f(p - m)$. ■

The maker quotes ask = $m + f$ and bid = $m - f$ and trades nothing in between: a proportional fee is a bid-ask spread, the same duality by which a proportional transaction cost confines the pricing functional to the bid-ask band (Jouini and Kallal 1995). Without convexity the exact description needs no conjugation at all:

Lemma 5 (the no-trade interval, exactly). *For any cost C (no convexity assumed), state q , and fee $f \geq 0$, write $d_q(s) = [C(q+s) - C(q)]/s$ for the chord slope. A belief μ admits no profitable trade if and only if*

$$\sup_{s < 0} d_q(s) - f \leq \mu \leq \inf_{s > 0} d_q(s) + f,$$

so the set of no-trade beliefs is exactly this interval intersected with the payoff hull. For differentiable convex C both bounds equal $C'(q)$ and the interval is Lemma 4's band $[m - f, m + f]$; at a convex kink the bounds are the one-sided derivatives, so the interval is $[\partial^- C(q) - f, \partial^+ C(q) + f]$ and the kink contributes a spread of its own even at $f = 0$. Write $\Delta_q = \sup_{s < 0} d_q - \inf_{s > 0} d_q$. The untruncated interval is non-empty if and only if $2f \geq \Delta_q$; if in addition C is chord-coherent, both bounds lie in the hull, so their midpoint does and the intersection with the hull is non-empty under the same condition. If C is chord-coherent and Proposition 2's supremum is attained, then off the contact set $\Delta_q > 0$.

Proof. No profitable trade means $\mu s \leq C(q+s) - C(q) + f|s|$ for all s . Dividing by $s > 0$ gives $\mu \leq d_q(s) + f$; dividing by $s < 0$ reverses the inequality to $\mu \geq d_q(s) - f$; together these are the displayed interval, whose non-emptiness is $\sup_{s < 0} d_q - f \leq \inf_{s > 0} d_q + f$. Under

chord coherence both one-sided bounds are limits of chord slopes and so lie in the hull, which is convex, so the midpoint of a non-empty interval lies in the hull. For the last claim, off the contact set Proposition 2 gives every belief a strictly positive frictionless profit, so the $f = 0$ interval is empty. ■

The hull intersection is not decoration. Without chord coherence $C(q) = 100q$ has $\Delta_q = 0$, so the untruncated interval $\{100\}$ is non-empty for every $f \geq 0$, yet no admissible belief in $[-1, 1]$ declines to trade: the maker is arbitrated from every state.

So friction does not merely widen an existing spread: at $2f \geq \Delta_q$ a state in the excluded inventory range becomes tenable, admitting no profitable trade at some admissible belief. Non-convexity makes an inventory range transient, frictionless rational flow jumps it (§3), and a large enough spread lets the maker rest inside it. For $C(q) = a \sin q$ at $q = 0$, a state sitting a above its flat envelope, the belief $\mu = 0$ admits no profitable trade exactly when $f \geq a$, though every belief profits there frictionlessly.

6 Makers in parallel

Let makers $i = 1..n$ hold inventories q_i with costs C_i that are differentiable and strictly convex, with C'_i onto the price range of interest so that $(C'_i)^{-1}$ is defined there; fees f_i of their choosing; and write $m_i = C'_i(q_i)$ for maker i 's marginal price and $\text{ask}_i = m_i + f_i$, $\text{bid}_i = m_i - f_i$ for its quotes, as in Lemma 4. Without differentiability and strict convexity the supply curves below are set-valued and the statements hold with $(C'_i)^{-1}$ read as a subgradient correspondence.

One normalization is needed before the merged object is a maker. The infimal convolution of the incremental costs need not vanish at zero: for two zero-fee quadratic makers $g_i(s) = m_i s + s^2/(2\lambda_i)$,

$$(g_1 \square g_2)(0) = -\frac{(m_1 - m_2)^2}{2(1/\lambda_1 + 1/\lambda_2)} < 0$$

whenever the makers' quotes differ, because zero net external demand still admits a profitable internal cross-trade. We therefore assume the component inventories have been cleared against each other and the effective cost normalized to vanish at the resulting origin; otherwise evaluating at $\Delta = 0$ repeatedly appears to pay the same internal arbitrage again and again.

Lemma 6 (combination). *Let $\tilde{C} = \tilde{C}_1 \square \dots \square \tilde{C}_n$ and define each maker's supply*

$$s_i(p) = \begin{cases} (C'_i)^{-1}(p - f_i) - q_i, & p \geq \text{ask}_i, \\ 0, & \text{bid}_i < p < \text{ask}_i, \\ (C'_i)^{-1}(p + f_i) - q_i, & p \leq \text{bid}_i, \end{cases}$$

each non-decreasing in p . Fix a demand Δ for which a clearing price p^ with $\sum_i s_i(p^*) = \Delta$ exists. Then: (i) $\tilde{C}^* = \sum_i \tilde{C}_i^*$, a sum of soft-thresholded profit functions; (ii) the split $s_i = s_i(p^*)$ attains $\tilde{C}(\Delta)$, and any optimal split satisfies $C'_i(q_i + s_i) + f_i \text{sign}(s_i) = p^*$ for $s_i \neq 0$ and $|p^* - m_i| \leq f_i$ for $s_i = 0$; (iii) the split is sparse: every maker whose quote band strictly contains p^* trades exactly zero.*

Proof. (i) is the conjugate-sum identity applied to the convolution (Rockafellar 1970). For (ii), the split is feasible by choice of p^* , and $p^* \in \partial \tilde{C}_i(s_i(p^*))$ for every i : when $s_i(p^*) \neq 0$ the subgradient is $C'_i(q_i + s_i) + f_i \text{sign}(s_i) = p^*$, and when $s_i(p^*) = 0$ it is the interval $[m_i -$

$f_i, m_i + f_i] \ni p^*$. A common multiplier certifying every coordinate is exactly the optimality condition for $\min\{\sum_i \tilde{C}_i(s_i) : \sum_i s_i = \Delta\}$. (iii) restates the zero branch. ■

The computation is a scalar monotone root-find whatever n is, and the fee costs nothing beyond a horizontal shift of each supply curve. The $|s|$ terms are an ℓ_1 penalty, so sparsity arrives for the same reason it does in the lasso, and for the same reason proportional transaction costs produce no-trade regions and sparse portfolios (Olivares-Nadal and DeMiguel 2018). A small trade routes to whichever maker posts the best quote, and to that maker alone unless the best quote is tied; a growing trade pushes that maker's fee-adjusted marginal price through the next band and spills over, consuming makers in quote-price order, which coincides with fee order only when their marginal prices agree.

Remark 7 (zero fees; known). With $f_i \equiv 0$ the convolution reduces to the fee-free merge: conjugate regularisers add, and for a perspective family $C_b(q) = b C_1(q/b)$ liquidity adds, $C_{b_1} \square C_{b_2} = C_{b_1+b_2}$. This is the aggregation law of Bhaskara et al. (2023), in risk-measure form Barrieu and El Karoui (2005), and is used here rather than claimed. For quadratic makers the merge is Gaussian fusion: the merged quote is the precision-weighted mean of the makers' quotes and the merged liquidity the sum of their precisions.

Corollary 8 (the order book). The aggregate supply $S(p) = \sum_i s_i(p)$ is non-decreasing, identically zero on $(\max_i \text{bid}_i, \min_i \text{ask}_i)$ when that interval is non-empty, flat on any open interval contained in the interiors of all the bands, and continuous and strictly increasing wherever some maker is active. If the active makers are in addition C^2 with $C_i'' > 0$ and p is away from band boundaries, the local depth is $S'(p) = \sum_{i \text{ active}} 1/C_i''(q_i + s_i(p))$. Differentiability needs that extra hypothesis: the maker $C(q) = q^4$ at $q_i = -1$ is differentiable and strictly convex with supply $S(p) = (p/4)^{1/3} + 1$, active at $p = 0$ and not differentiable there. Read as a market: best bid and ask are the tightest quotes, $S(p)$ is the cumulative quantity executable up to price p with local depth $S'(p)$ where it exists, and large orders walk the levels. The aggregate of linear-fee makers is a consolidated limit order book, and in producer-theory terms Lemma 6 is Marshall's horizontal summation of firm supply curves (Marshall 1890; Mas-Colell et al. 1995) with the reversibility of share production patched by the fee. The economics of the book assembled from competing liquidity suppliers is Glosten (1994), with convergence of strategic schedules in Biais et al. (2000).

Proposition 9 (a deep co-quoter is Moreau smoothing). Merging a maker with cost C with a quadratic co-quoter of cost $s^2/(2\lambda)$ yields the venue with cost the Moreau envelope $e_\lambda C(x) = \min_y C(y) + (x - y)^2/(2\lambda)$. Then: (i) coherence is preserved: if the chord slopes of C lie in $[a, b]$, so do those of $e_\lambda C$, for every λ ; (ii) at any point where two distinct minimizing branches coexist (a crossing), the downward jump in marginal price is exactly the separation of the competing minimizers divided by λ , at most D/λ when all competing minimizers lie in an interval of diameter D ; (iii) for the symmetric double well $C(y) = \alpha \min(|y - c|, |y + c|)$ with $\alpha < 1$, the jump at $x = 0$ equals 2α while $\lambda\alpha < c$ and $2c/\lambda$ thereafter, so it is constant in the depth until the co-quoter is deep enough to reach across the wells. The bound in (ii) measures crossing defects only: the envelope can be smoothly non-convex with a unique minimizer everywhere, as for $C(y) = a \sin y$ with $\lambda a < 1$, where $(e_\lambda C)''(x) = C''(y^*)/(1 + \lambda C''(y^*))$ is negative wherever C'' is.

Proof. The merge is infimal convolution, and inf-convolution with the quadratic is the Moreau envelope (Rockafellar 1970). For (i), no differentiability is needed: for $h > 0$, substituting $y = z + h$,

$$e_\lambda C(x+h) = \inf_z \left\{ C(z+h) + \frac{(x-z)^2}{2\lambda} \right\} \leq e_\lambda C(x) + bh,$$

since $C(z+h) - C(z) \leq bh$ for every z ; the reverse substitution gives the lower bound ah , so every chord slope of $e_\lambda C$ lies in $[a, b]$. For (ii), at a crossing x_0 with competing minimizers $y_1^* < y_2^*$ the branch slopes are $(x_0 - y_i^*)/\lambda$, so the drop is exactly $(y_2^* - y_1^*)/\lambda$. For the smooth case, differentiate the first-order condition $y^* = x - \lambda C'(y^*)$ to get $dy^*/dx = 1/(1 + \lambda C''(y^*))$ and hence the displayed curvature. For (iii), the well's competing minimizers sit at $\pm\lambda\alpha$ while $\lambda\alpha < c$, giving separation $2\lambda\alpha$ and jump 2α , and at $\pm c$ once the co-quoter reaches the wells, giving $2c/\lambda$. ■

Part (i) is the durable statement: adding depth reshapes a non-convex maker without ever creating arbitrage, at any λ . Depth does not, however, fill the excluded range at finite depth. For the double well $e_\lambda C(\pm c) = 0$ while $e_\lambda C(0) = \alpha c - \lambda\alpha^2/2$ for $\lambda\alpha < c$ and $c^2/(2\lambda)$ after, strictly positive for every finite λ ; since $e_\lambda C \geq 0$ vanishes at $\pm c$, its convex envelope is zero at the midpoint, so the midpoint stays strictly off contact at every depth and the gap merely decays like $1/\lambda$. The accurate statement is that depth attenuates the envelope gap and regularizes crossing geometry without creating chord incoherence; it does not close the off-contact interval, restore convexity, or make every state reachable. Crossing kinks close at rate $1/\lambda$ only where the competing minimizers stay within a λ -independent diameter; the double well shows what happens otherwise, its jump holding at 2α until the depth exceeds c/α . Smooth concave stretches lie outside the bound's scope entirely. Here λ is the co-quoter's liquidity: its price impact is $1/\lambda$, so large λ means a deep book, not a strong pull.

The two repairs are the two market primitives again: friction (ℓ_1 , the fee) and participant depth (ℓ_2 , capital), lasso and ridge. The identification maps the penalty terms, not the statistical procedures; it is a dictionary of primitives, not an equivalence of estimators.

7 Self-set fees and adaptivity

Nothing requires the fees to be administered. Each maker may quote its own f_i , and routing then disciplines the posted menu: a maker quoting inside the prevailing spread becomes the best quote and takes the flow first, a maker quoting behind the best price receives nothing until the book in front of it is consumed, and a maker quoting too tight is picked off by informed flow.

The discipline is a property of Lemma 6's routing, not of its minimization. The infimal convolution minimizes over allocations s_i with the fees held fixed; it does not minimize over the fees, and nothing in it constitutes a game in which makers choose f_i . What the surviving fee would be, and whether it is the competitive adverse-selection charge of classical microstructure (Glosten and Milgrom 1985; Biais et al. 2000), requires the strategic quote game, which is not solved here.

Optimization already pays these frictions: a proximal step charges $\|\Delta\theta\|^2/(2\eta)$ per move, a trust region is an infinite fee outside a band, weight decay is a zero-quoting participant. In market language the stabilizers of non-convex training are the market's two repairs: friction prices chord excursions (Proposition 3) and depth attenuates the envelope gap (Proposition 9); non-convexity by itself is not exploitability (§2). Path-dependent optimizers correspond to makers whose quotes depend on flow history, the adaptive-liquidity territory where path independence is deliberately traded away: no maker combines path independence, translation

invariance (complete-bundle prices summing to the sure payoff) and liquidity sensitivity (Othman et al. 2013), the adaptive class is axiomatized by Li and Wortman Vaughan (2013) with the homogeneous-risk-measure characterization in Othman and Sandholm (2011), no trade-history maker achieves every desideratum at once (Abernethy et al. 2014), and liquidity selection itself can be run as online learning (Nueve et al. 2026; Nueve and Waggoner 2025).

8 Inconsistency is arbitrage

Sections 2–7 concern one venue and its aggregate. The same convex geometry says when *several* venues, quoting related securities, are jointly incoherent, and it turns each incoherence into a portfolio that pays whoever finds it.

Proposition 10 (quote inconsistency is a sure profit). *Let x be a vector of unrestricted real random variables and let separate venues quote securities paying the products $x_i x_j$ for every pair, the diagonal quoted at one, at fixed linear prices executable in the size traded; assemble the quotes into a matrix P with unit diagonal. The quotes admit a joint distribution iff $P \succeq 0$ (a Gaussian with covariance P witnesses sufficiency). Otherwise let w be a unit eigenvector of a negative eigenvalue $\lambda_{\min}$ and buy w_i^2 units of security ii and $2w_i w_j$ units of each security ij , $i < j$. The bundle costs $w^\top P w = \lambda_{\min} < 0$ and pays $(w^\top x)^2 \geq 0$: a sure profit of $|\lambda_{\min}|$. On three coordinates the quoted pairs are exactly the edges of a triangle, so the loopy reading is literal.*

This is the second-moment coherence condition, prices of products must form a PSD matrix (d’Aspremont 2005), read as cycle consistency. Arbitrage-enforced coherence is an old idea. Coherence is no-arbitrage (Nau and McCardle 1991); Pennock and Wellman (1996) built arbitrageur agents enforcing the additivity identities of a Bayes-net economy, with equilibrium prices equal to the network’s probabilities and distributed bidding as distributed inference; the combinatorial literature removes the arbitrage algorithmically over the marginal polytope (Kroer et al. 2016); and Sagui et al. (2025) measure Polymarket arbitrageurs extracting \$39.6M enforcing logical coherence across markets.

Two boundaries of the proposition mark the general picture. If only a subgraph of pairs is quoted, the realizability condition is not positive semidefiniteness of a filled-in matrix but PSD-completability of the partial one (Grone et al. 1984), and the arbitrage bundle must be supported on quoted pairs. If the coordinates are constrained, say binary $x_i \in \{\pm 1\}$, the feasible quote set is the cut polytope, a strict subset of the PSD body (Deza and Laurent 1997), so quotes can be positive semidefinite yet jointly unrealizable. In every case the coherent set is the convex hull of attainable payoffs, and Proposition 11 below is the universal form.

The scope of Proposition 10 is in any case one class of consistency failure: locally quoted beliefs that cannot be embedded in any joint distribution, detected here at second moments. Proposition 10 is a second-moment arbitrage result, not a theorem about loopy belief propagation in general: not every loopy-propagation error takes this form, and the reading of arbitrageurs as the loop correction is a conjecture about dynamics that this paper does not model. What it does establish is that the class it covers is, in a market, free money. Whether flow that harvests it converges, and to the true marginal, to a surrogate of the kind loopy propagation converges to, or to something the liquidity profile selects, is open.

The PSD cone is not special. Coherent price sets are convex (they are convex hulls of payoff vectors, or their conic images), and separation turns every exterior point into a trade:

Proposition 11 (arbitrage is separation). *Let K be the closed convex hull of the attainable payoff vectors (or its image under a linear security map), so that $\inf_{z \in K} \langle y, z \rangle \leq \langle y, \varphi(\omega) \rangle$ for every outcome*

ω and portfolio y , and let the venue quote fixed linear prices x , executable in the size traded. If $x \notin K$, any separating functional y with $\langle y, x \rangle < \inf_{z \in K} \langle y, z \rangle$ is a portfolio whose price is strictly less than its realized payoff in every state: a sure profit. Proposition 10 is the instance $K = \{P \succeq 0\}$.

For a nonlinear cost-function maker the conclusion is local. There x is the current gradient and a trade of size δy costs $\delta \langle y, x \rangle + o(\delta)$, since the marginal price moves as the trade is filled, so a separating margin exceeding $f \|y\|$ yields sure profit for all sufficiently small δ , at a rate rather than in the size of y ; a margin below that is suppressed by Proposition 3, and market impact then sets how large a position the certificate supports.

The proof is the separating-hyperplane theorem read as a trade (Nau and McCardle 1991; d’Aspremont 2005); the payoff-hull hypothesis is what upgrades the separation certificate to an arbitrage, and for an abstract consistency set the construction yields the certificate only. In an ordinary numerical method an infeasibility is a residual to be driven down by the algorithm; in a market it is a payoff, and whoever finds it is paid to act as the separation oracle. Arbitrageurs are decentralized separation oracles, and the friction of §4 sets the tolerance below which infeasibility is allowed to persist.

9 Closing

The characterization is for the cost-based class. Call a predictor *cost-based* if it is specified by a path-independent potential C over a security inventory, and call friction of size f *permitted* if the mechanism may charge $f \|s\|$ on a fill s . For vector inventories arbitrage depth is measured against the same norm,

$$\varepsilon = \sup_{q, s \neq 0} \frac{[\inf_{z \in K} \langle z, s \rangle - C(q + s) + C(q)]_+}{\|s\|},$$

the scalar case of §4 being $\|\cdot\| = |\cdot|$. A cost-based predictor admits a coherent market implementation if and only if its arbitrage depth is dominated by the permitted friction. The equivalence is close to definitional, arbitrage depth being the maximal normalized violation, and it is included to fix vocabulary rather than to carry weight; what does carry weight is that the depth is finite for costs no axiom set in the literature admits, and that §§2–4 identify what finiteness costs. Chord coherence is the $f = 0$ special case, and convexity is the separate, further property of monotone information incorporation. Three words are kept apart throughout: *coherence* (chord slopes stay in the payoff hull), *convexity* (marginal prices respond monotonically to flow), and *expressiveness* (which inventory states rational flow can leave the maker in). Within the proper-envelope class, and for beliefs whose response is well posed, non-convexity restricts the rationally reachable inventory states to the contact set: C and $\hat{C}$ give the same optimal value, with

$$\operatorname{argmax}_x \{\mu x - C(x)\} = \operatorname{argmax}_x \{\mu x - \hat{C}(x)\} \cap \{C = \hat{C}\},$$

so the original cost deletes the envelope’s off-contact optimizers, and from an off-contact state the next rational trader extracts the gap $g(q)$. Coherence alone promises no more than this, since $C(q) = -\alpha|q|$ is coherent with unbounded speculative profit (§3).

The target operation throughout is *marketization*, of which this paper realizes the single-maker and parallel parts; the serial part is the subject of the companion paper. Given an op-

erator T from inputs to outputs, a marketization is a mechanism whose clearing computes T ; whose local contributions compose, in parallel and in series, to compute composite operators; in which inconsistent contributions create exploitable trades (Proposition 11); in which friction bounds how much inconsistency can persist (Proposition 3); and in which every participant pays to perturb the computation. One estimator becomes one maker and combining evidence becomes parallel composition; composing conditional operators is serial composition, developed separately. A statistical procedure specifies how information would be combined if supplied honestly; its marketization implements the same operator against self-interested sources. In this sense a market is the incentive closure of a predictor: the computation, plus the dual certificates of its constraints, plus payment to whoever holds one. Duality alone supplies the middle column of

computation	certificate	payment
primal solve	dual multiplier	none
verified computation	proof of violation	none or fixed
incentive closure	separating portfolio	the certificate's value

and the closure adds the third: a marketized computation is self-policing in proportion to the value of its own errors, with friction the exchange rate between tolerated error and the cost of participation.

The definition has neighbors in four literatures, and each holds one column. Markets that compute are old: a combinatorial auction cleared airport slots by integer programming (Rassenti et al. 1982), electricity markets pay the dual variables of the dispatch program as nodal prices (Schweppe et al. 1988), and market-oriented programming solved distributed allocation by computing competitive equilibria of artificial economies (Wellman 1993). Algorithmic mechanism design asks when optimization is compatible with incentives (Nisan and Ronen 2001), and the classification of what can be computed incentive-compatibly was posed as a program by Feigenbaum and Shenker (2002). The middle column carries a lower bound: any communication protocol realizing an efficient allocation must reveal supporting prices (Nisan and Segal 2006), so certificates are not optional. The third column exists in practice without the duality: fraud-proof systems pay whoever exhibits an invalid state transition (Deutsch and Reitwießner 2019; Kalodner et al. 2018), for discrete transitions rather than convex programs. And the friction bound has an incentive-free ancestor: the auction algorithm's ε -complementary slackness permits suboptimality at most $n\varepsilon$, vanishing for $\varepsilon < 1/n$ on integer data (Bertsekas 1992); Proposition 3 is the same shape of bound with the ε charged to traders. The connective tissue, the closure operation itself, serial composition of mechanisms, and the identification of violation certificates with separating portfolios, appears unoccupied.

One principle deserves separation from the open problems it generates, because it qualifies everything above.

Principle (accessible arbitrage). *Market coherence is relative not only to friction but to the cost of discovering violations. Write K for the exactly coherent quotes (Proposition 1) and K_f for those admitting no sure profit net of a fee f (Proposition 3). Let $\mathcal{A}$ be a class of admissible certificate-search procedures containing the zero-cost null procedure, let x encode the venue's executable books and depth rather than a vector of marginal quotes, and for a portfolio π let*

$$V_f(\pi; x) = \inf_{\omega} [\text{payoff}_{\omega}(\pi) - \text{executionCost}_x(\pi) - \text{fees}_f(\pi)]$$

be the guaranteed profit, maximized over the positions the arbitrageur's depth, capital and position

limits allow, with $C_A(x)$ the search cost and the expectation below taken over the randomness of the search procedure alone. The accessible-coherence set is

$$K_{\mathcal{A},f} = \left\{ x : \sup_{A \in \mathcal{A}} \mathbb{E}[V_f(A(x); x) - C_A(x)] \leq 0 \right\}.$$

With nonnegative search costs, $K \subseteq K_f \subseteq K_{\mathcal{A},f}$: a violation persists whenever every available separation procedure costs at least as much, net of friction, as the certificate it finds is worth.

Every feature of V_f is load-bearing. The supremum ranges over procedures, not certificates: two procedures can find the same portfolio at very different costs, so coherence is relative to a class of arbitrageurs. The null procedure makes the supremum at least zero, so that $K \subseteq K_f \subseteq K_{\mathcal{A},f}$ holds. The profit is a worst case over outcomes, not realized or belief-weighted profit, without which exact coherence would not imply that every procedure has non-positive value. And the profit must be realizable. A separating functional may be scaled by any positive constant, so against fixed linear prices its nominal value is unbounded and no finite search cost could ever leave a violation standing; it is the nonlinear book, the finite executable depth, and the position constraint that make $K_f \subsetneq K_{\mathcal{A},f}$ possible at all. Equivalently one may normalize certificates to $\|y\| \leq 1$ and price them against available depth. Instantiating $\mathcal{A}$ gives polynomial-time coherence, bounded-budget coherence, and latency-constrained coherence, and a venue can be arbitrage-free to ordinary participants while arbitrageable to a better-equipped class. Three nested notions of coherence result: exact (no sure-profit chord, Proposition 1), frictional (violations below the spread survive, Proposition 3), and accessible (violations too expensive to discover survive). Proposition 11 addresses the existence of arbitrage; whether a procedure in $\mathcal{A}$ can profitably prove $x \notin K$ is its accessibility, a different property. The self-policing of the marketization table is qualified accordingly: effective self-correction is the error's value net of discovery cost and friction.

Open problems, in rough order of tractability:

Market representations of general prediction maps. The characterization above is confined to cost-based predictors. Define what it means for an arbitrary prediction map T , from data sets to forecasts, to possess a market representation, and give conditions on T equivalent to existence. A first conjecture: T is representable when it factors through a network of cost kernels whose boundary reduction is T , with the chord bound applying kernel by kernel; path independence and the chord condition would then be the image of the general condition in the cost-based case.

Accessible coherence, characterized. Compute $K_{\mathcal{A},f}$ for concrete classes: polynomial-time arbitrageurs against cones whose separation is hard, budgeted arbitrageurs against fee-bearing venues, latency-constrained arbitrageurs against cascading settlements. And characterize the mechanisms in which the stream of certificate payments funds the ongoing computation.

The equilibrium fee. Section 7 argues discipline, not equilibrium. In a flow model where most coefficients are zero and informed trades carry signal plus noise, does the fee that survives undercutting reproduce the universal threshold $\sigma(2 \log p)^{1/2}$ of the shrinkage literature? If so, the regularization constant that statistics tunes by cross-validation equals the adverse-selection cost of the data source.

Arbitrage depth as a capacity measure. On a specified compact domain, with a fixed payoff hull and norm, the arbitrage depth of a trained model's loss landscape is well defined and plausibly estimable or boundable, though computing it exactly is a global problem; whether it tracks quantities learning theory names (sharpness, mode connectivity) is open.

From endpoint jumps to quoted depth. Section 3 predicts an inventory jump at a single supporting price, and notes the jump is generally lumpy rather than divisible. The map from that endpoint correspondence to anything observable in a quoted book is missing, and is what a test against venues with non-convex liquidity, or against the bimodal implied densities documented around binary events (Melick and Thomas 1997; Clark and Amen 2017), would require first.

References

- Abernethy, Jacob, Yiling Chen, and Jennifer Wortman Vaughan. 2013. “Efficient Market Making via Convex Optimization, and a Connection to Online Learning.” *ACM Transactions on Economics and Computation* 1 (2): 12:1–39. <https://doi.org/10.1145/2465769.2465777>.
- Abernethy, Jacob, Rafael Frongillo, Xiaolong Li, and Jennifer Wortman Vaughan. 2014. “A General Volume-Parameterized Market Making Framework.” *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 413–30. <https://doi.org/10.1145/2600057.2602900>.
- Angeris, Guillermo, Tarun Chitra, Theo Diamandis, Alex Evans, and Kshitij Kulkarni. 2024. “The Geometry of Constant Function Market Makers.” *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, 732. <https://doi.org/10.1145/3670865.3673500>.
- Barbu, Adrian, and Nathan Lay. 2012. “An Introduction to Artificial Prediction Markets for Classification.” *Journal of Machine Learning Research* 13: 2177–204. <https://arxiv.org/abs/1102.1465>.
- Barrieu, Pauline, and Nicole El Karoui. 2005. “Inf-Convolution of Risk Measures and Optimal Risk Transfer.” *Finance and Stochastics* 9 (2): 269–98. <https://doi.org/10.1007/s00780-005-0152-0>.
- Bertsekas, Dimitri P. 1992. “Auction Algorithms for Network Flow Problems: A Tutorial Introduction.” *Computational Optimization and Applications* 1 (1): 7–66. <https://doi.org/10.1007/BF00247653>.
- Beygelzimer, Alina, John Langford, and David M. Pennock. 2012. “Learning Performance of Prediction Markets with Kelly Bettors.” *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. <https://arxiv.org/abs/1201.6655>.
- Bhaskara, Adithya, Rafael Frongillo, Elias Lindgren, and Maneesha Papireddygar. 2023. *A General Theory of Liquidity Provisioning for Prediction Markets*. arXiv:2311.08725. <https://arxiv.org/abs/2311.08725>.
- Biais, Bruno, David Martimort, and Jean-Charles Rochet. 2000. “Competing Mechanisms in a Common Value Environment.” *Econometrica* 68 (4): 799–837. <https://doi.org/10.1111/1468-0262.00138>.

- Chen, Yiling, and Jennifer Wortman Vaughan. 2010. "A New Understanding of Prediction Markets via No-Regret Learning." *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, 189–98. <https://doi.org/10.1145/1807342.1807372>.
- Clark, Iain J., and Saeed Amen. 2017. "Implied Distributions from GBPUSD Risk-Reversals and Implication for Brexit Scenarios." *Risks* 5 (3): 35. <https://doi.org/10.3390/risks5030035>.
- Cummings, Rachel, David M. Pennock, and Jennifer Wortman Vaughan. 2016. "The Possibilities and Limitations of Private Prediction Markets." *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. <https://arxiv.org/abs/1602.07362>.
- d'Aspremont, Alexandre. 2005. *A Market Test for the Positivity of Arrow-Debreu Prices*. arXiv:cs/0510027. <https://arxiv.org/abs/cs/0510027>.
- Deza, Michel M., and Monique Laurent. 1997. *Geometry of Cuts and Metrics*. Vol. 15. Algorithms and Combinatorics. Springer. <https://doi.org/10.1007/978-3-642-04295-9>.
- El Ghaoui, Laurent, and Hervé Lebrete. 1997. "Robust Solutions to Least-Squares Problems with Uncertain Data." *SIAM Journal on Matrix Analysis and Applications* 18 (4): 1035–64. <https://doi.org/10.1137/S0895479896298130>.
- Feigenbaum, Joan, and Scott Shenker. 2002. "Distributed Algorithmic Mechanism Design: Recent Results and Future Directions." *Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, 1–13. <https://www.cs.yale.edu/homes/jf/FS.pdf>.
- Frongillo, Rafael M., Nicolás Della Penna, and Mark D. Reid. 2012. "Interpreting Prediction Markets: A Stochastic Approach." *Advances in Neural Information Processing Systems* 25 (NeurIPS). <https://papers.nips.cc/paper/4529-interpreting-prediction-markets-a-stochastic-approach>.
- Frongillo, Rafael, Maneesha Papireddygar, and Bo Waggoner. 2024. "An Axiomatic Characterization of CFMMs and Equivalence to Prediction Markets." *15th Innovations in Theoretical Computer Science Conference (ITCS)*, LIPIcs, vol. 287: 51:1–21. <https://doi.org/10.4230/LIPIcs.ITCS.2024.51>.
- Frongillo, Rafael, and Bo Waggoner. 2018. "Bounded-Loss Private Prediction Markets." *Advances in Neural Information Processing Systems (NeurIPS)*. <https://arxiv.org/abs/1703.00899>.
- Glosten, Lawrence R. 1994. "Is the Electronic Open Limit Order Book Inevitable?" *The Journal of Finance* 49 (4): 1127–61. <https://doi.org/10.1111/j.1540-6261.1994.tb02450.x>.
- Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics* 14 (1): 71–100.

- Grone, Robert, Charles R. Johnson, Eduardo M. Sá, and Henry Wolkowicz. 1984. "Positive Definite Completions of Partial Hermitian Matrices." *Linear Algebra and Its Applications* 58: 109–24. [https://doi.org/10.1016/0024-3795\(84\)90207-6](https://doi.org/10.1016/0024-3795(84)90207-6).
- Guasoni, Paolo. 2006. "No Arbitrage Under Transaction Costs, with Fractional Brownian Motion and Beyond." *Mathematical Finance* 16 (3): 569–82. <https://doi.org/10.1111/j.1467-9965.2006.00283.x>.
- Guasoni, Paolo, Miklós Rásonyi, and Walter Schachermayer. 2010. "The Fundamental Theorem of Asset Pricing for Continuous Processes Under Small Transaction Costs." *Annals of Finance* 6 (2): 157–91. <https://doi.org/10.1007/s10436-008-0110-x>.
- Hanson, Robin. 2003. "Combinatorial Information Market Design." *Information Systems Frontiers* 5 (1): 107–19. <https://doi.org/10.1023/A:1022058209073>.
- Jouini, Elyès, and Hédi Kallal. 1995. "Martingales and Arbitrage in Securities Markets with Transaction Costs." *Journal of Economic Theory* 66 (1): 178–97. <https://doi.org/10.1006/jeth.1995.1037>.
- Kalodner, Harry, Steven Goldfeder, Xiaoqi Chen, S. Matthew Weinberg, and Edward W. Felten. 2018. "Arbitrum: Scalable, Private Smart Contracts." *Proceedings of the 27th USENIX Security Symposium*, 1353–70. <https://www.usenix.org/conference/usenixsecurity18/presentation/kalodner>.
- Kroer, Christian, Miroslav Dudík, Sébastien Lahaie, and Sivaraman Balakrishnan. 2016. "Arbitrage-Free Combinatorial Market Making via Integer Programming." *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 161–78. <https://arxiv.org/abs/1606.02825>.
- Lépinette, Emmanuel, and Tuan Quoc Tran. 2017. "Arbitrage Theory for Non Convex Financial Market Models." *Stochastic Processes and Their Applications* 127 (10): 3331–53. <https://www.sciencedirect.com/science/article/pii/S0304414915301216>.
- Li, Xiaolong, and Jennifer Wortman Vaughan. 2013. "An Axiomatic Characterization of Adaptive-Liquidity Market Makers." *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, 657–74. <https://doi.org/10.1145/2482540.2482575>.
- Marshall, Alfred. 1890. *Principles of Economics*. Macmillan.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. Oxford University Press.
- Melick, William R., and Charles P. Thomas. 1997. "Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil During the Gulf Crisis." *Journal of Financial and Quantitative Analysis* 32 (1): 91–115. <https://doi.org/10.2307/2331318>.

- Nau, Robert F., and Kevin F. McCardle. 1991. "Arbitrage, Rationality, and Equilibrium." *Theory and Decision* 31 (2–3): 199–240.
- Nisan, Noam, and Amir Ronen. 2001. "Algorithmic Mechanism Design." *Games and Economic Behavior* 35 (1–2): 166–96. <https://doi.org/10.1006/game.1999.0790>.
- Nisan, Noam, and Ilya Segal. 2006. "The Communication Requirements of Efficient Allocations and Supporting Prices." *Journal of Economic Theory* 129 (1): 192–224. <https://doi.org/10.1016/j.jet.2004.10.007>.
- Nueve, Enrique, Bao Nguyen, Rafael Frongillo, and Bo Waggoner. 2026. *Adaptive Liquidity in Prediction Markets via Online Learning*. arXiv:2605.09599. <https://arxiv.org/abs/2605.09599>.
- Nueve, Enrique, and Bo Waggoner. 2025. *Smooth Quadratic Prediction Markets*. arXiv:2505.02959. <https://arxiv.org/abs/2505.02959>.
- Olivares-Nadal, Alba V., and Victor DeMiguel. 2018. "A Robust Perspective on Transaction Costs in Portfolio Optimization." *Operations Research* 66 (3). <https://doi.org/10.1287/opre.2017.1699>.
- Othman, Abraham, David M. Pennock, Daniel M. Reeves, and Tuomas Sandholm. 2013. "A Practical Liquidity-Sensitive Automated Market Maker." *ACM Transactions on Economics and Computation* 1 (3). <https://doi.org/10.1145/2509413.2509414>.
- Othman, Abraham, and Tuomas Sandholm. 2011. "Liquidity-Sensitive Automated Market Makers via Homogeneous Risk Measures." *Internet and Network Economics (WINE 2011)*, Lecture notes in computer science, vol. 7090. https://doi.org/10.1007/978-3-642-25510-6_27.
- Othman, Abraham, and Tuomas Sandholm. 2012. "Profit-Charging Market Makers with Bounded Loss, Vanishing Bid/Ask Spreads, and Unlimited Market Depth." *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, 790–807. <https://doi.org/10.1145/2229012.2229074>.
- Pennock, David M., and Michael P. Wellman. 1996. "Toward a Market Model for Bayesian Inference." *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence (UAI)*, 405–13. <https://arxiv.org/abs/1302.3593>.
- Pennock, David M., and Michael P. Wellman. 1997. "Representing Aggregate Belief Through the Competitive Equilibrium of a Securities Market." *Proceedings of the 13th Conference on Uncertainty in Artificial Intelligence (UAI)*, 392–400. <https://arxiv.org/abs/1302.1564>.
- Rassenti, Stephen J., Vernon L. Smith, and Robert L. Bulfin. 1982. "A Combinatorial Auction Mechanism for Airport Time Slot Allocation." *The Bell Journal of Economics* 13 (2): 402–17. <https://doi.org/10.2307/3003463>.
- Rockafellar, R. Tyrrell. 1970. *Convex Analysis*. Princeton Mathematical Series 28. Princeton

- University Press. <https://doi.org/10.1515/9781400873173>.
- Saguillo, Oriol, Vahid Ghafouri, Lucianna Kiffer, and Guillermo Suarez-Tangil. 2025. "Unraveling the Probabilistic Forest: Arbitrage in Prediction Markets." *Advances in Financial Technologies (AFT)*. <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.AFT.2025.27>.
- Schweppe, Fred C., Michael C. Caramanis, Richard D. Tabors, and Roger E. Bohn. 1988. *Spot Pricing of Electricity*. Kluwer Academic Publishers.
- Storkey, Amos J. 2011. "Machine Learning Markets." *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, PMLR, vol. 15: 716–24. <https://arxiv.org/abs/1106.4509>.
- Storkey, Amos J., Jono Millin, and Krzysztof Geras. 2012. "Isoelastic Agents and Wealth Updates in Machine Learning Markets." *International Conference on Machine Learning (ICML)*. <https://arxiv.org/abs/1206.6443>.
- Teutsch, Jason, and Christian Reitwießner. 2019. *A Scalable Verification Solution for Blockchains*. arXiv:1908.04756. <https://arxiv.org/abs/1908.04756>.
- Wellman, Michael P. 1993. "A Market-Oriented Programming Environment and Its Application to Distributed Multicommodity Flow Problems." *Journal of Artificial Intelligence Research* 1: 1–23. <https://doi.org/10.1613/jair.2>.
- Xu, Huan, Constantine Caramanis, and Shie Mannor. 2009. "Robustness and Regularization of Support Vector Machines." *Journal of Machine Learning Research* 10: 1485–510. <https://jmlr.org/papers/v10/xu09b.html>.