An Algebra of Prediction-Rewarding Mechanisms

Scoring rules, market makers, and pools as composable transducers

Peter Cotton

Working draft v0.4 · 2026

Abstract

Scoring rules, cost-function market makers, constant-function market makers, parimutuel pools, and opinion pools are developed in separate literatures. This note is expository: it gathers the connections among them under one convex-analytic lens and one notational convention. A proper scoring rule is a convex potential read off the report; the Fenchel conjugate of that potential is a cost-function market maker, and a level-set dual is a constant-function market maker. The linear and logarithmic opinion pools are the two Kullback-Leibler barycenters. Merging market makers is infimal convolution, the composition law of the risk-sharing literature, so liquidity adds. A common transducer signature lets the mechanisms be chained, and this note settles when propriety survives a transformation of a stage's message or outcome. The mathematics is classical; the contribution is its consolidation. The multi-stage theory built on this dictionary, and its deployment on the microprediction platform, is a companion report (Cotton 2026b); two further results that motivated the framing are elsewhere too, sample-based elicitation on point clouds (Cotton 2026c) and the parimutuel account of a conformal predictor's information gap (Cotton 2026a).

1 Introduction

Mechanisms for eliciting and aggregating forecasts are usually studied one at a time. A proper scoring rule is analysed as a one-shot contract; a market maker as a sequential trading venue; an opinion pool as an estimator; a calibration test as a diagnostic. Deployed systems are thinner than the theory allows: with few exceptions they run one level of competition against one internal model. Numerai pools staked submissions into a single stake-weighted meta-model and pays each forecast its marginal contribution to it (Craib et al. 2017); CrunchDAO blends each contest into one ensemble; the IARPA prediction polls aggregated forecasts by track record, comparing favourably with head-to-head markets (Atanasov et al. 2017), with reputation rather than wealth as the threaded state. In every case, one pool and one internal aggregate.

Studied in separate literatures, these mechanisms look unrelated. Under convex duality they are one object seen from different sides: a proper scoring rule, a cost-function market maker, a constant-function market maker, and the two opinion pools are all read off a single convex potential and its conjugate. Writing that dictionary down, in one convention, is the purpose of this note. The mechanisms also chain, one stage's belief output becoming the next stage's message or settlement, and the multi-stage theory together with its deployment on the microprediction platform is a companion report (Cotton 2026b); here the object is the single-stage dictionary the chain is built from.

This note organizes the single-stage dictionary around three questions that are often blurred together:

1. *Message closure.* Can every stage consume and emit one common object, so that stages plug together at all?
2. *Convex generation.* Is each individual mechanism generated by a convex potential, so that one dictionary covers scoring rules, market makers, and their duals?
3. *Propriety under transformation.* When a stage transforms forecasts or outcomes, does strict propriety survive?

The first, message closure, is the organizing thesis of the companion report on multi-stage solicitation (Cotton 2026b); this note takes it as a convention (§2) and settles the second and third, which admit theorems, stated with complete proofs in §3 and §4. The mathematics is classical; the purpose here is to collect it in one place and one convention and to record which connections hold. Two developments that motivated the framing are carried elsewhere and not reproved: sample-based elicitation on the finite clouds deployed contests collect (Cotton 2026c), and the parimutuel account of a conformal predictor's information gap (Cotton 2026a), summarized as an instance in §5.

2 Preliminaries

The finite setting. Section 3 works with a finite outcome set $\{1, \ldots, n\}$ and reports p in the probability simplex Δ; from the aggregation operators of §4 onward the paper moves to the continuous setting, densities and distribution functions on $\mathbb{R}$ or $[0, 1]^d$. The finite statements transfer with the usual measure-theoretic care; the one place where the transfer is not routine is the probability integral transform, treated separately in Proposition 4.

Scores. A scoring rule assigns $S(p, i) \in \mathbb{R} \cup \{-\infty\}$ to a report p and outcome i; the expected score of reporting p under belief q is $S(p; q) = \sum_i q_i S(p, i)$, assumed well defined with $S(q; q)$ finite (this is the regularity used throughout). The rule is *proper* if $S(q; q) \geq S(p; q)$ for all p, q and *strictly proper* if equality forces $p = q$. Scores are written in reward form: higher is better.

Convex tools. For convex G with subgradient selection $G'(p) \in \partial G(p)$, the Bregman divergence is $D_G(q, p) = G(q) - G(p) - \langle G'(p), q - p \rangle$. The Legendre-Fenchel conjugate of a function R on Δ is $R^*(\mathbf{q}) = \sup_{p \in \Delta} \langle p, \mathbf{q} \rangle - R(p)$. The infimal convolution of f and g is $(f \,\square\, g)(x) = \inf_y f(y) + g(x - y)$. "Closed proper convex" is used in the standard sense (Rockafellar 1970). Differentiable statements are on the relative interior of Δ, with extended-real values allowed on the boundary; subgradients on Δ act on its tangent space and are defined up to addition of multiples of $\mathbf{1}$.

3 One potential, three mechanisms

Everything in this section is one fact in different clothing: a scoring rule is proper exactly when it is built from a convex potential. The honest forecaster's expected score is that potential, and the penalty for any other report is the potential's curvature between truth and report. The three classical scores are three choices of potential, and the market makers of the next results are its conjugate.

Theorem 1 (characterisation of proper scoring rules; Savage (1971), McCarthy (1956)). *A regular scoring rule S is proper iff there is a convex $G : \Delta \to \mathbb{R}$ with*

$$S(p,i) \;=\; G(p) + \langle\, G'(p),\; e_i - p\,\rangle, \qquad G'(p) \in \partial G(p) \quad (p \in \operatorname{ri}\Delta),$$

where e_i is the i-th unit vector. It is strictly proper iff G is strictly convex relative to Δ, and then $G(p) = S(p;p)$ is the expected score of a truthful forecaster.

The idea is that the expected score is affine in the belief, so the gain from honesty $S(q;q) - S(p;q)$ collapses to the Bregman divergence $D_G(q,p)$, which a supporting hyperplane makes non-negative and strict convexity makes zero only at $q = p$.

Proof. ($\Leftarrow$) Since $\sum_i q_i(e_i - p) = q - p$, the expected score is affine in the belief: $S(p;q) = G(p) + \langle G'(p), q - p\rangle$. Hence

$$S(q;q) - S(p;q) \;=\; G(q) - G(p) - \langle G'(p),\, q - p\rangle \;=\; D_G(q,p),$$

which is non-negative by the supporting-hyperplane inequality, and zero only at $q = p$ when G is strictly convex relative to Δ; so S is (strictly) proper.

($\Rightarrow$) Define $G(q) := S(q;q)$. Properness says $G(q) = \sup_p S(p;q)$, and each $q \mapsto S(p;q) = \langle S(p,\cdot), q\rangle$ is affine, so G is a pointwise supremum of affine functions, hence convex. For any fixed p,

$$S(p;q) \;=\; G(p) + \langle S(p,\cdot),\, q - p\rangle \quad \text{with} \quad S(p;q) \le G(q)\;\; \forall q,$$

so the vector $S(p,\cdot)$ is a subgradient of G at p (acting on the tangent space of Δ, where subgradients are defined up to a constant shift along $\mathbf{1}$, which the representation absorbs). Substituting $G'(p) = S(p,\cdot)$ into the display returns $S(p,i)$ identically. For strictness, note $D_G(q,p) = 0$ exactly when G has a supporting hyperplane at p that also supports G at q. If G carried a nontrivial affine segment in Δ, a hyperplane supporting it at an interior point of the segment would support it along the whole segment, giving $D_G(q,p) = 0$ for a distinct $q \neq p$, so S would not be strictly proper; conversely strict convexity relative to Δ forces $D_G(q,p) = 0$ only at $q = p$. ■

The content of Theorem 1 is the dictionary *proper scoring rule ↔ convex function ↔ Bregman divergence* (Gneiting and Raftery 2007; Banerjee et al. 2005). The classics are three choices of G, with the divergences computed directly from the definition:

Score	Generator $G(p) = S(p;p)$	Bregman divergence $D_G(q,p)$
logarithmic	$\sum_i p_i \log p_i$	$\mathrm{KL}(q\\|p)$
Brier (quadratic)	$\\|p\\|_2^2$	$\\|q-p\\|_2^2$
spherical	$\\|p\\|_2$	$\\|q\\|_2\,(1-\cos\theta_{p,q})$

For the spherical row: $\nabla G(p) = p/\|p\|$, so $D_G(q,p) = \|q\| - \langle p,q\rangle/\|p\| = \|q\|(1-\cos\theta_{p,q})$, an angular term scaled by $\|q\|$. For the logarithmic row the representation gives $S(p,i) = \log p_i$ on the relative interior, with $-\infty$ on the boundary.

The same potential runs a market. A cost-function market maker posts a running cost C for bundles of outcome shares and lets traders move the state; reading the generator G of Theorem 1 as a regulariser and taking its convex conjugate produces such a C, and the market's prices, its freedom from arbitrage, and its worst-case loss are then three properties of that conjugate. Hanson's LMSR is the entropic case.

Theorem 2 (scoring rule to market maker; Hanson (2007), Abernethy et al. (2013)). *Let R be closed, proper, strictly convex, and finite-valued on Δ (in particular $R(e_i) < \infty$ at every vertex, which the worst-case-loss bound needs), extended by $+\infty$ off Δ so that $C = R^*$ below is the Fenchel conjugate on $\mathbb{R}^n$ (in the sequel, $R = G$, the generator of Theorem 1 read as a regulariser), and define*

$$C(\mathbf{q}) = \sup_{p \in \Delta} \big(\langle p, \mathbf{q}\rangle - R(p)\big).$$

Then:

(i) C is convex and finite, the maximiser $p^*(\mathbf{q})$ is unique, and $\nabla C(\mathbf{q}) = p^*(\mathbf{q}) \in \Delta$: prices are a probability vector. (Without strict convexity, prices live in $\partial C(\mathbf{q})$.)

(ii) C is translation-equivariant, $C(\mathbf{q} + \alpha\mathbf{1}) = C(\mathbf{q}) + \alpha$; costs telescope over any trade path, so round trips cost zero and buying the full bundle costs its payout: the market is arbitrage-free.

(iii) With initial state $\mathbf{q}_0 = \mathbf{0}$, the maker's loss when the market settles on outcome i after net sales $\mathbf{q}$ is

$$\mathrm{loss}_i(\mathbf{q}) = q_i - C(\mathbf{q}) + C(\mathbf{0}) \;\le\; R(e_i) - \inf_{p\in\Delta} R(p) \;\le\; \sup_{p\in\Delta} R(p) - \inf_{p\in\Delta} R(p).$$

(iv) Taking $R(p) = b\sum_i p_i \log p_i$ gives $C(\mathbf{q}) = b\log\sum_i e^{q_i/b}$, Hanson's LMSR, with worst-case loss $b\log n$.

Each part is a standard reading of the conjugate: prices are its gradient (Danskin), no-arbitrage is its translation-equivariance, and the loss bound is the range of the regulariser (Fenchel-Moreau).

Proof. (i) C is a supremum of affine functions of $\mathbf{q}$, hence convex; the supremum of a continuous function over the compact Δ is attained, and strict concavity of $p \mapsto \langle p, \mathbf{q}\rangle - R(p)$ makes the maximiser unique. Danskin's theorem gives $\nabla C(\mathbf{q}) = p^*(\mathbf{q})$. (ii) $\langle p, \mathbf{q} + \alpha\mathbf{1}\rangle = \langle p, \mathbf{q}\rangle + \alpha$ for $p \in \Delta$, so the supremum shifts by α. A trader moving the state $\mathbf{q} \to \mathbf{q}'$ pays $C(\mathbf{q}') - C(\mathbf{q})$ by definition, so costs over any path telescope; a round trip costs zero, and by translation equivariance the bundle $\alpha\mathbf{1}$ costs exactly α, its sure payout. (iii) The maker collects $C(\mathbf{q}) - C(\mathbf{0})$ and pays q_i, so $\mathrm{loss}_i = q_i - C(\mathbf{q}) + C(\mathbf{0})$. Then $\sup_{\mathbf{q}} \big(\langle e_i, \mathbf{q}\rangle - C(\mathbf{q})\big) = C^*(e_i) = R^{**}(e_i) = R(e_i)$ (Fenchel-Moreau, R closed), and $C(\mathbf{0}) = \sup_p -R(p) = -\inf_p R(p)$; hence the supremum of the loss over $\mathbf{q}$ is exactly $R(e_i) - \inf_p R(p)$. The final inequality holds because $e_i \in \Delta$. (iv) Lagrange: maximising $\langle p, \mathbf{q}\rangle - b\sum p_i \log p_i$ subject to $\sum p_i = 1$ gives $q_i - b(\log p_i + 1) = \lambda$, so $p_i \propto e^{q_i/b}$, and substituting back yields $C(\mathbf{q}) = b\log\sum_i e^{q_i/b}$. Then $R(e_i) = 0$ and $\inf R = -b\log n$ at the uniform distribution, so the loss bound is $b\log n$. ■

Proposition 3 (cost-function markets and CFMMs). *Under the monotonicity, concavity, and reserve-domain hypotheses of Frongillo et al. (2024), cost-function prediction markets and constant-function market makers can be converted into one another. In the unconstrained-reserve sign convention a cost function C gives a concave CFMM potential by $\varphi(\mathbf{r}) = -C(-\mathbf{r})$; bounded-reserve versions require a perspective (level-set) construction. The equivalence is a convex level-set duality; it is not the bare Fenchel conjugacy $C \mapsto C^*$.*

We do not reprove the general equivalence here; Frongillo et al. (2024) give it in full, and Angeris et al. (2023) and Angeris et al. (2021) develop the CFMM side.

Example (constant product). Take two outcome tokens with reserves r_1, r_2 and invariant $r_1 r_2 = k$. The pool's portfolio value at prices $(p, 1-p)$ is

$$V(p) = \inf\{p\, r_1 + (1-p)\, r_2 : \, r_1 r_2 = k\} = 2\sqrt{k\, p(1-p)},$$

by the AM-GM inequality, with the infimum attained at $r_1 = \sqrt{k(1-p)/p}$, $r_2 = \sqrt{kp/(1-p)}$. V is concave; reading $R(p) = -V(p)$ as the regulariser of Theorem 2 gives a maker whose worst-case loss is the range of R on $[0,1]$, namely $\sqrt{k}$ (between $p = \frac{1}{2}$ and the endpoints): the geometric mean of the initial reserves. The corresponding generator is not entropic, which is the convex-analytic content of the observation that constant-product markets and the LMSR price bounded-payout claims differently.

Proposition 4 (the probability integral transform; Rosenblatt (1952), Dawid (1984)). *If X has continuous CDF F then $U = F(X) \sim \mathrm{Uniform}(0,1)$, and $z = \Phi^{-1}(U) \sim N(0,1)$. In the prequential setting, if F_t is the true conditional law of X_t given the past, then $U_t = F_t(X_t)$ is conditionally uniform and the PIT stream is iid uniform.*

Proof. For continuous F, with the generalized inverse $F^{-}(u) = \inf\{x : F(x) \geq u\}$, the event $\{F(X) \leq u\}$ differs from $\{X \leq F^{-}(u)\}$ by a set of probability zero, and $\Pr(X \leq F^{-}(u)) = F(F^{-}(u)) = u$ by continuity. The prequential statement applies this conditionally at each t. ■

Two scope remarks. First, the converse fails in the strong sense: marginally uniform PITs do not certify an informative forecast. Reporting the unconditional law F of an iid sequence gives $F(X_t) \sim \mathrm{Uniform}(0,1)$ even if valuable covariates were ignored. A PIT critic therefore witnesses miscalibration relative to its test class; it complements, and does not replace, a proper score that rewards sharp conditional distributions. Conformal prediction lives on exactly this distinction: a split-conformal predictor achieves marginal coverage by construction, the uniform-PIT guarantee, while free to ignore the conditional information in the input; §5 recounts the gap (Cotton 2026a). Second, for discrete forecasts the randomized PIT preserves the exact uniform null, while the mid-PIT is a convenient deterministic diagnostic with a different null distribution.

4 Operators on mechanisms

The common signature. A *transducer* (Mealy machine) is a tuple $(S, A, B, \delta, \lambda)$ with state space S, input alphabet A, output alphabet B, transition map $\delta : S \times A \to S$, and output map $\lambda : S \times A \to B$; run on an input stream it produces $s_{t+1} = \delta(s_t, a_t)$ and $b_t = \lambda(s_t, a_t)$, a causal map of input streams to output streams. The mechanisms of §3 share one instantiation. Let Dist denote the set of distributional beliefs and take S the wealth states, $A = \mathrm{Dist}^m \times \mathcal{X}$ (m participant reports and, where the stage settles, a realized outcome), and $B = \mathrm{Dist} \times \mathbb{R}^m$ (an aggregate belief and transfers). A *stage* is such a transducer, written

$$M : (\mathrm{Dist}^m,\, w,\, x) \longmapsto (\mathrm{Dist},\, w',\, \pi).$$

A scoring rule is the transfer component λ of a stage, not itself a map $\mathrm{Dist} \to \mathrm{Dist}$; a market maker is a stage whose state is the inventory vector and whose emitted belief is the price; an opinion pool is a stage with no outcome argument and zero transfers. Composition wires the belief output of one stage to the belief inputs of the next while state and transfers thread through.

The operators below act on stages.

Sequentialise. Theorem 2: a proper score run sequentially against a wealth state is a cost-function market maker.

Pool. A proper score gives a batch elicitation mechanism when reports are scored independently and funded externally: properness is inherited report by report. Parimutuel and budget-balanced versions are a different game, because the pot split couples payoffs through the denominator; the price-taking analysis (Cotton 2026c, sec. 2) gives truthful all-in, a symmetric equilibrium at fractional stakes, and degeneracy as the stake fraction vanishes, and beyond that the equilibrium theory is open.

Averaging beliefs has two natural senses, one for each direction of the Kullback-Leibler divergence, and they are exactly the two pools used in practice: the linear pool is the barycenter that pulls the average toward each forecast, the logarithmic pool the one that pulls each forecast toward the average.

Ensemble (Proposition 5: the two pools are the two KL barycenters). *Let $q_1, \dots, q_m$ be densities with respect to a common dominating measure, $w_i \geq 0$, $\sum w_i = 1$, and for the second display assume $0 < \int \prod_i q_i^{w_i} < \infty$. Then the linear pool minimizes the forward divergences and the logarithmic pool the reverse:*

$$\arg\min_p \sum_i w_i \operatorname{KL}(q_i \| p) = \sum_i w_i q_i, \qquad \arg\min_p \sum_i w_i \operatorname{KL}(p \| q_i) \ \propto\ \prod_i q_i^{w_i}.$$

Proof. For the first, $\sum_i w_i \operatorname{KL}(q_i \| p) = \text{const} - \int (\sum_i w_i q_i) \log p$, and $\int \bar{q} \log p$ is maximized over densities at $p = \bar{q}$ (Gibbs). For the second, $\sum_i w_i \operatorname{KL}(p \| q_i) = \int p \log p - \int p \overline{\log q}$ with $\overline{\log q} = \sum_i w_i \log q_i$; the Lagrange condition is $\log p = \overline{\log q} + \text{const}$. ∎

Training weights by cumulative log score and then mixing linearly, $\sum_m w_m F_m$, is Bayesian model averaging, a linear pool with score-trained weights; the logarithmic pool multiplies densities and renormalizes. They are different aggregates with different sharpness (Genest and Zidek 1986).

Two market makers standing back to back act as one deeper maker, and the depth adds because conjugation turns the sum of their regularisers into the infimal convolution of their cost functions.

Merge (Proposition 6: merging makers is infimal convolution; Rockafellar (1970), Bhaskara et al. (2023)). *For closed proper convex f, g the conjugate identity $(f \,\square\, g)^* = f^* + g^*$ holds unconditionally; recovering the primal form $f \,\square\, g = (f^* + g^*)^*$ needs $f \,\square\, g$ closed and the infimal convolution exact (automatic for the finite-valued regularisers here, since $\operatorname{ri}\operatorname{dom} R_1 \cap \operatorname{ri}\operatorname{dom} R_2 = \operatorname{ri}\Delta \neq \emptyset$). Consequently, merging two cost-function makers with regularisers R_1, R_2 (cost functions $C_i = R_i^*$) yields the maker with regulariser $R_1 + R_2$, and merging LMSR_{b_1} with LMSR_{b_2} yields $\text{LMSR}_{b_1+b_2}$: liquidity adds.*

Proof. $(f \,\square\, g)^*(p) = \sup_x \langle p, x\rangle - \inf_y \{f(y) + g(x-y)\} = \sup_{y,z} \langle p, y\rangle - f(y) + \langle p, z\rangle - g(z) = f^*(p) + g^*(p)$. For the makers, $C_1 \,\square\, C_2 = (R_1 + R_2)^*$ by biconjugacy, and $b_1 G + b_2 G = (b_1 + b_2) G$ for the entropic generator. ∎

This is also the risk-sharing literature's composition law: the aggregate of several agents' convex risk measures is their infimal convolution, with the Pareto allocation as minimiser and the common subgradient as the clearing price (Barrieu and El Karoui 2005; Jouini et al. 2008); the market-liquidity reading is developed by Bhaskara et al. (2023).

Conjugate. Run a stage in a transformed coordinate and map back. This is where propriety needs care, and the theory splits in two:

- *Fixed transformations.* Properness is preserved under any fixed transformation of forecast and outcome, and strict propriety iff the transformation is injective (Allen et al. 2023, Prop. 4; Pic et al. 2025, Prop. 1). The Markov-kernel (stochastic channel) extension, strict propriety iff the channel is injective on laws, is Theorem 2 of the companion point-cloud paper, whose Theorem 1 also exhibits the canonical failure: a KDE smoothing seam whose raw-outcome score elicits a deconvolution.
- *Forecast-dependent transformations.* The PIT critic transforms outcomes by the reported F itself, so the fixed-transformation theorems do not apply to it; Proposition 4 and its scope remarks are the correct warrant, and the critic is a calibration diagnostic rather than a strictly proper score. The stacked-lottery design (Cotton 2020, slides 29-31) is this operator in practice: percentiles from one game feed the next, and calibration is produced by composing monotone maps contributed by competing algorithms.

5 A conformal predictor as a degenerate composition

Split-conformal prediction is a composition with its middle stage switched off: a point predictor feeds a rank-based calibration stage, but the pool step is skipped and all credit is assigned to one model in advance. The calibration stage prices the residual flat in the input, so its marginal coverage is exact while any conditional structure in the residual is left unpriced. Run that residual stage as an actual parimutuel and the gap is collected: an entrant who conditions on the input X grows their bankroll at the rate $I(R; X)$, the mutual information between residual and input, which is exactly the conditional information the flat predictor discards. Marginal coverage is the break-even statement; the conditional information is the rent.

The full account, with the pool payoff, the Kelly-Breiman growth identity, the Gaussian rate $-\frac{1}{2}\log(1-\rho^2)$, and the anytime-valid measurement, is a standalone note (Cotton 2026a); the mechanism ran in production as the microprediction nearest-the-pin pool and the MidOne residual contest, reported in the companion paper on multi-stage solicitation (Cotton 2026b).

6 Implementation notes

The mechanisms of the catalogue are implemented against a single interface patterned on the skaters time-series contract (Cotton 2026d), a successor to timemachines (Cotton 2021): every stage consumes and emits the same distributional type and threads its own state, so the signature of §4 is the forecasting contract with wealth in place of model state. The closure of the message type is what makes a small operator set sufficient; the transforms, ensembles, and residual constructions of §4 all consume and emit the one type. Scores are implemented in loss form; the text uses reward form.

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.

Allen, Sam, David Ginsbourger, and Johanna Ziegel. 2023. "Evaluating Forecasts for High-Impact Events Using Transformed Kernel Scores." *SIAM/ASA Journal on Uncertainty Quantification* 11 (3): 906–40. https://doi.org/10.1137/22M1532184.

Angeris, Guillermo, Tarun Chitra, Alex Evans, and Matthew Lorig. 2023. "Short Communication: A Primer on Perpetuals." *SIAM Journal on Financial Mathematics* 14 (1): SC17–30. https://doi.org/10.1137/22M1520931.

Angeris, Guillermo, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. 2021. "An Analysis of Uniswap Markets." *Cryptoeconomic Systems* 1 (1). https://doi.org/10.21428/58320208.c9738e64.

Atanasov, Pavel, Phillip Rescober, Eric Stone, et al. 2017. "Distilling the Wisdom of Crowds: Prediction Markets Vs. Prediction Polls." *Management Science* 63 (3): 691–706. https://doi.org/10.1287/mnsc.2015.2374.

Banerjee, Arindam, Xin Guo, and Hui Wang. 2005. "On the Optimality of Conditional Expectation as a Bregman Predictor." *IEEE Transactions on Information Theory* 51 (7): 2664–69. https://doi.org/10.1109/TIT.2005.850145.

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.

Bhaskara, Adithya, Rafael Frongillo, and Maneesha Papireddygari. 2023. *A General Theory of Liquidity Provisioning for Prediction Markets*. arXiv:2311.08725. https://arxiv.org/abs/2311.08725.

Cotton, Peter. 2020. *The Lottery Paradox: A New Use*. https://www.slideshare.net/slideshow/lottery-paradox-csaildec2020/242173597.

Cotton, Peter. 2021. *Timemachines: Time-Series Prediction with a Uniform Skater Interface*. https://github.com/microprediction/timemachines.

Cotton, Peter. 2026a. *Betting Against a Conformal Predictor: A Parimutuel Account of the Information Gap*. https://conformalprediction.net/papers/parimutuel/.

Cotton, Peter. 2026b. *Multi-Stage Solicitation of Probability Distributions: Experiments, Theory and Perspective on Conformal Prediction*. https://mechanisms.microprediction.org/papers/multi-stage-solicitation.html.

Cotton, Peter. 2026c. *Scoring Point-Cloud Distributional Submissions*. https://mechanisms.microprediction.org/papers/scoring-point-cloud-distributional-submissions.html.

Cotton, Peter. 2026d. *Skaters: Fast Online Time-Series Models with Composable Transforms*. https://github.com/microprediction/skaters.

Craib, Richard, Geoffrey Bradway, Xander Dunn, and Joey Krug. 2017. *Numeraire: A Cryptographic Token for Coordinating Machine Intelligence*. https://numer.ai/whitepaper.pdf.

Dawid, A. P. 1984. "Present Position and Potential Developments: Some Personal Views. Statistical Theory: The Prequential Approach." *Journal of the Royal Statistical Society. Series A (General)* 147 (2): 278–92. https://doi.org/10.2307/2981683.

Frongillo, Rafael, Maneesha Papireddygari, 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.

Genest, Christian, and James V. Zidek. 1986. "Combining Probability Distributions: A Critique and an Annotated Bibliography." *Statistical Science* 1 (1): 114–35. https://doi.org/10.1214/ss/1177013825.

Gneiting, Tilmann, and Adrian E. Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association* 102 (477): 359–78. https://doi.org/10.1198/016214506000001437.

Hanson, Robin. 2007. "Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation." *The Journal of Prediction Markets* 1 (1): 3–15. https://doi.org/10.5750/jpm.v1i1.417.

Jouini, Elyes, Walter Schachermayer, and Nizar Touzi. 2008. "Optimal Risk Sharing for Law Invariant Monetary Utility Functions." *Mathematical Finance* 18 (2): 269–92. https://doi.org/10.1111/j.1467-9965.2007.00332.x.

McCarthy, John. 1956. "Measures of the Value of Information." *Proceedings of the National Academy of Sciences* 42 (9): 654–55. https://doi.org/10.1073/pnas.42.9.654.

Pic, Romain, Clément Dombry, Philippe Naveau, and Maxime Taillardat. 2025. "Proper Scoring Rules for Multivariate Probabilistic Forecasts Based on Aggregation and Transformation." *Advances in Statistical Climatology, Meteorology and Oceanography* 11 (1): 23–58. https://doi.org/10.5194/ascmo-11-23-2025.

Rockafellar, R. Tyrrell. 1970. *Convex Analysis*. Princeton Mathematical Series 28. Princeton University Press. https://doi.org/10.1515/9781400873173.

Rosenblatt, Murray. 1952. "Remarks on a Multivariate Transformation." *The Annals of Mathematical Statistics* 23 (3): 470–72. https://doi.org/10.1214/aoms/1177729394.

Savage, Leonard J. 1971. "Elicitation of Personal Probabilities and Expectations." *Journal of the American Statistical Association* 66 (336): 783–801. https://doi.org/10.1080/01621459.1971.10482346.