Clamped Markets

How prediction mechanisms compose, and what familiar methods are when the competition is switched off

Peter Cotton

Working draft v0.1 · August 29, 2026

Abstract

This paper is partly synthetic. Many of the correspondences below are known individually; the purpose is to place them in one compositional calculus and to ask what a familiar statistical method is when read through it. The organizing observation is that many familiar statistical methods are market ecosystems with degrees of freedom clamped. Mechanisms compose in two ways, and both are established. In parallel, competing makers merge by infimal convolution, so conjugates add and liquidity adds. In series, a factorization assigns one venue per conditional factor, the arrangement Hanson's modularity makes tradable and that deployed combinatorial markets price by junction-tree inference. Reading the two laws through the generalized distributive law puts them in one semiring frame, in which a market of optimizing traders computes min-plus and agrees with Bayesian inference on the log-quadratic family. Against that background, clamping is the operation that recovers ordinary statistics: freeze entry and a market of Kelly bettors becomes Bayesian model averaging; freeze capital and a market of data points becomes least squares; freeze the propagation step and a chain of venues becomes the Kalman filter; freeze the residual venue to a single unconditional entry and a two-stage solicitation becomes split-conformal prediction. Each clamp forfeits a specific rent, and in the conformal case the rent is exactly the conditional information $I(R;X)$ that the unconditional entry declines to price. The paper sets out the composition laws with their sources, catalogues the clamps, and states what unclamping would require. Where the calculus stops is itself informative: it locates the mechanism that is missing.

1 Two composition laws, and a dial

A prediction mechanism takes beliefs and pays for them. Two ways of putting mechanisms together account for most of what is known.

Parallel. Several makers quote the same quantity. Their aggregate is the infimal convolution of their cost functions, so conjugates add and liquidity adds; this is the aggregation law of Bhaskara et al. (2023), the risk-sharing composition of Barrieu and El Karoui (2005) and Jouini et al. (2008), and in the cost-function framework of Abernethy et al. (2013) it is the statement that merging two makers yields a deeper one. For quadratic makers it is Gaussian fusion: the merged quote is the precision-weighted mean and precisions add.

Serial. A model is a factorization, $P(\text{everything}) = \prod_i P(x_i \mid \text{parents})$, and the serial arrangement runs one venue per factor, each pricing its conditional given what is upstream. The locality that makes this tradable is Hanson's modularity: in a combinatorial logarithmic market scoring rule a bet on $A \mid B$ moves that conditional and provably nothing else, uniquely among market scoring rules (Hanson 2007, 2003).

Neither law is new here, and §§2–3 set them out with their sources. What this paper adds is a dial. Between a full ecosystem, in which every stage is a live venue with free entry, and a single fitted model, there is a continuum of *clamped* mechanisms: the same composition, with some degree of freedom frozen. Ordinary statistical procedures sit at the clamped end, and the question worth asking of each is not whether it is a market but which clamp it applies and what that clamp costs. Section 5 catalogues them, and conformal prediction turns out to be the sharpest case, because there the forfeited rent has a closed form.

Three kinds of statement appear below and are worth distinguishing as they arrive. *Known identities* are used and cited, not claimed: that trading against a cost-function maker is follow-the-regularized-leader, that a market of Kelly bettors performs Bayesian model averaging, that merging makers is infimal convolution, that Gaussian fusion is information-form addition, that min-plus elimination is dynamic programming. *Results* are the few small formal statements, mostly in the companion paper, that the calculus turns up on its way through. And *principles* are proposed organizing ideas that are not yet theorems, chiefly clamping itself and the accounting of what each clamp forfeits. The individual identities are mostly known; the algebra connecting them, and what it says about methods that are not usually thought of as markets, is the subject here.

2 Parallel composition: what is settled

Cost-function market makers price by a convex potential, with prices its gradient, no-arbitrage from convexity and bounded loss from a bounded conjugate range (Abernethy et al. 2013; Hanson 2003). The duality with proper scoring rules is standard: the maker is the conjugate of the scoring rule's entropy, and trading against it is follow-the-regularized-leader (Chen and Vaughan 2010).

Merging is infimal convolution. Bhaskara et al. (2023) prove that liquidity providers submitting arbitrary cost functions operate, in parallel, exactly as the maker whose cost is $\bigwedge_i C_i$, with the dual generating functions adding and the split across providers computed behind the scenes; Angeris et al. (2024) read Minkowski sums of trade sets the same way. In risk-measure language this is the classical result that the aggregate of convex risk measures is their infimal convolution, with the Pareto allocation as minimizer and the common subgradient as clearing price (Barrieu and El Karoui 2005; Jouini et al. 2008).

Fees complicate the picture and are treated in the companion paper: a proportional fee is exactly a bid-ask spread, participation becomes sparse, and the aggregate supply curve is a consolidated limit order book. Nothing here depends on that development.

3 Serial composition: what is settled

Pricing a combinatorial market *is* probabilistic inference, which is why exact pricing of an LMSR over a combinatorial space is #P-hard (Chen, Fortnow, et al. 2008); why tournament markets price by Bayes-net inference (Chen, Goel, et al. 2008); why a deployed combinatorial market ran its price and asset updates on the junction-tree algorithm (Sun et al. 2012); and why tractable designs price approximately over the marginal polytope (Dudík et al. 2012; Dudík et al. 2021).

Which securities preserve a Bayes-net structure under trade is characterized by Xia and Pennock (2011), and securities structured by a factorization appear in Pennock and Wellman (2000).

That trading and message passing are the same activity is likewise established. Pennock and Wellman (1996) map a Bayes net to an economy with one agent per conditional-probability entry and arbitrageur agents enforcing the additivity identities, showing equilibrium prices equal the network's probabilities and distributed bidding is distributed inference. Storkey (2011) shows that agents caring about subsets of variables give equilibria factorizing as products of local potentials, and derives messages from optimized positions.

The arrangement called serial here differs in architecture rather than in that observation: separate venues with settlement boundaries and an opening rule, each factor a market that opens when its conditioning information freezes and settles in cascade, rather than one joint book priced by an inference algorithm. Whether self-interested trading can be made to perform the elimination itself, rather than a solver performing it inside one venue, is open, and is the subject of the companion paper's final section.

4 One semiring

The two laws are the two operations of a commutative semiring, which is the frame the generalized distributive law provides (Aji and McEliece 2000): one operation combines evidence about a variable, one moves evidence between variables, and sum-product, max-product and min-sum are one algorithm over different semirings. Assigning each factor its potential $\varphi_i = -\log p_i$, parallel merge adds potentials (equivalently inf-convolves costs, the two related by the Legendre transform, which is the min-plus Laplace transform (Litvinov 2005)), and serial composition is the min-plus kernel product $\inf_y[\varphi_1(x, y) + \varphi_2(y, z)]$.

Two consequences are worth recording because they are often run together. First, a market of optimizing traders computes min-plus natively: a finite-state chain of stage potentials prices exactly the Viterbi decoding of the corresponding hidden Markov model, with no Gaussian structure anywhere. Second, min-plus agrees with sum-product on log-quadratic families, where partial minimization equals marginalization up to a constant independent of the retained variables, by the Schur-complement identity. On that family the alternation of a propagation step with a fusion step reproduces the Kalman filter (Kalman 1960) exactly, the fusion being a parallel merge, and messages on a Gaussian chain reproduce belief propagation (Pearl 1988). Off it, the two diverge by the Laplace-approximation gap, and what a market prices is the max-marginal rather than the marginal.

5 Clamping

Now the dial. In each case below a mechanism of §§2–4 is recovered from a familiar procedure by *unfreezing* something, and the procedure is recovered by clamping it again.

Freeze entry: Bayesian model averaging. A market of Kelly bettors prices the wealth-weighted average of participants' beliefs, and wealth updates are Bayes updates, so the market performs Bayesian model averaging with wealths as posterior weights (Beygelzimer et al. 2012; Kelly 1956; Breiman 1961). Model averaging over a fixed model list is that market with entry closed. What the clamp forfeits is whatever a model outside the list would have earned; the market-selection literature is the study of who survives when entry is open (Blume and Easley 2006).

Freeze capital: least squares. 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, the estimate their merge. Fixing those capitals at the nominal precisions is a clamp: the market version lets a source that has been wrong lose capital, which is the difference between weighted least squares and its robust or adaptive variants.

Freeze the fee: frictionless aggregation. A maker with no fee earns nothing on uninformed flow and cannot recover adverse-selection losses from volume. Setting fees to zero is the clamp that makes an aggregation rule out of a venue; letting participants quote their own fee is what turns the spread into a price. The companion paper develops this.

Freeze propagation: the Kalman filter. In the filtering recursion the correction step is a parallel merge, executed by trades against an observation maker, while the prediction step is computed by whoever owns the model. That asymmetry is a clamp on the edge factor: nobody is paid to move the state forward. Unclamping it means posting the edge as a venue with securities on the pair (x_t, x_{t+1}), which is the open problem the companion paper states.

Freeze the residual venue: conformal prediction. This is the sharpest case and the one that motivates the vocabulary.

6 Conformal prediction as a clamped two-stage market

Split-conformal prediction forms a residual law from calibration data and applies it at every input. Read as a mechanism it is a two-stage solicitation: a base predictor, then a residual stage. The second stage is a market with the field restricted to a single *unconditional* entry (Cotton 2026b).

The distinction that matters is not which contest a forecaster enters but whether their submission moves with the covariates. An unconditional forecaster submits one distribution and stands by it; a conditional forecaster submits one that depends on the input she sees. In a log-wealth pool settling on an outcome, the best conditional entrant out-earns the best unconditional one by exactly the conditional information the covariates carry, $I(R;X)$, the mutual information between the residual and the input (Cotton 2026a). That number is the rent the clamp forfeits, and in the Gaussian case it is $-\frac{1}{2}\log(1-\rho^2)$.

Two consequences follow. Marginal coverage, the guarantee conformal prediction offers (Lei et al. 2018), is the break-even statement of the clamped venue: the unconditional entry is priced correctly on average, which is what makes it safe and also what makes it leave money on the table. And the gap does not close when the conformists are right: if the base forecast is calibrated so the residuals are marginally standard, a conditional participant still profits at the same rate, because marginal correctness is not conditional correctness (Foygel Barber et al. 2021). Adaptive conformal variants that let the law depend on the input are conditional entries, which is to say they have partially unclamped the venue and are forecasting.

Run the residual stage as an actual pool and the rent is collected rather than forfeited. That mechanism ran in production as the microprediction platform's `z1~` residual streams and its MidOne contest (Cotton 2026b).

7 Point-cloud games, and the unclamped end

At the other end of the dial sits a mechanism with nothing frozen. In a nearest-the-pin parimutuel over a continuum, participants submit clouds of samples, the pot is split in proportion to the density each placed at the realized outcome, and entry is open; this was the reward engine behind montepredicition (Cotton 2024). Every degree of freedom the clamped methods freeze is live here: who enters, how much each stakes, what each conditions on, and how sharp a submission is.

Point-cloud submission also shows what the unclamped end costs. Scoring a kernel density at the raw outcome elicits the *deconvolution* of the participant's belief rather than the belief, an impropriety with a closed form, repaired by jittering the pin with the same kernel; and in high dimensions the pool runs on random projections. Those are the subject of a separate paper, and the point here is only that an open mechanism has design problems a clamped one never encounters, because a clamped mechanism has already decided the answers.

The platform stacked such games: a pool on a live quantity, a stream predicting that pool's own calibration, dependence streams pricing copulas, and a lottery of calibration maps (Cotton 2026b). Read through this calculus, that architecture is the serial law of §3 with every stage left unclamped, and the familiar methods of §5 are what remain when the stages are frozen one at a time.

8 What clamping costs, and what unclamping would require

The catalogue suggests a common shape. Each clamp replaces a price with an assumption:

clamp	assumption replacing a price	forfeited
entry closed	the model list is adequate	what an outside model would earn
capital fixed	nominal precisions are right	the correction from realized performance
fee zero	flow is uninformed	adverse-selection revenue
propagation computed	the dynamics are known	the price of the transition
residual unconditional	covariates carry nothing	$I(R;X)$

Only the last entry currently has a closed form, and finding the others is the natural program. Two further problems are worth naming. Unclamping a stage requires a settlement rule for that stage's securities, and a latent quantity with no settlement functional is not a market however convenient it is to speak of one. And the serial law, though it is what the platform implemented, has no general theorem saying self-interested trading performs the elimination; the companion paper states what such a mechanism would need.

The claim is modest and, I think, useful: the individual identities are mostly known, and the algebra connecting them is the subject here. The clamped limits of these mechanisms are the methods statistics already uses. Naming the clamp is a quick way to see what a method has given up, and in one case it says exactly how much.

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.

Aji, Srinivas M., and Robert J. McEliece. 2000. "The Generalized Distributive Law." *IEEE Transactions on Information Theory* 46 (2): 325–43. https://doi.org/10.1109/18.825794.

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.

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.

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 Papireddygari. 2023. *A General Theory of Liquidity Provisioning for Prediction Markets*. arXiv:2311.08725. https://arxiv.org/abs/2311.08725.

Blume, Lawrence, and David Easley. 2006. "If You're so Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets." *Econometrica* 74 (4): 929–66. https://doi.org/10.1111/j.1468-0262.2006.00691.x.

Breiman, Leo. 1961. "Optimal Gambling Systems for Favorable Games." *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* 1: 65–78.

Chen, Yiling, Lance Fortnow, Nicolas Lambert, David M. Pennock, and Jennifer Wortman. 2008. "Complexity of Combinatorial Market Makers." *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, 190–99. https://doi.org/10.1145/1386790.1386822.

Chen, Yiling, Sharad Goel, and David M. Pennock. 2008. "Pricing Combinatorial Markets for Tournaments." *Proceedings of the 40th ACM Symposium on Theory of Computing (STOC)*, 305–14. https://doi.org/10.1145/1374376.1374421.

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.

Cotton, Peter. 2024. *Monteprediction: A Weekly Monte Carlo Competition in Eleven Dimensions*. https://www.monteprediction.com.

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/mult i-stage-solicitation.html.

Dudík, Miroslav, Xintong Wang, David M. Pennock, and David M. Rothschild. 2021. "Log-Time Prediction Markets for Interval Securities." *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 465–73. https://doi.org/10.65109/bbzi6501.

Dudík, Miroslav, Sébastien Lahaie, and David M. Pennock. 2012. "A Tractable Combinatorial Market Maker Using Constraint Generation." *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, 459–76. https://doi.org/10.1145/2229012.2229047.

Foygel Barber, Rina, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. 2021. "The Limits of Distribution-Free Conditional Predictive Inference." *Information and Inference* 10 (2): 455–82. https://doi.org/10.1093/imaiai/iaaa017.

Hanson, Robin. 2003. "Combinatorial Information Market Design." *Information Systems Frontiers* 5 (1): 107–19. https://doi.org/10.1023/A:1022058209073.

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.

Kalman, Rudolf E. 1960. "A New Approach to Linear Filtering and Prediction Problems." *Journal of Basic Engineering* 82 (1): 35–45. https://doi.org/10.1115/1.3662552.

Kelly, Jr., J. L. 1956. "A New Interpretation of Information Rate." *The Bell System Technical Journal* 35 (4): 917–26. https://doi.org/10.1002/j.1538-7305.1956.tb03809.x.

Lei, Jing, Max G'Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. 2018. "Distribution-Free Predictive Inference for Regression." *Journal of the American Statistical Association* 113 (523): 1094–111. https://doi.org/10.1080/01621459.2017.1307116.

Litvinov, Grigory L. 2005. *The Maslov Dequantization, Idempotent and Tropical Mathematics: A Brief Introduction*. arXiv:math/0507014. https://arxiv.org/abs/math/0507014.

Pearl, Judea. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.

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. 2000. "Compact Securities Markets for Pareto Optimal Reallocation of Risk." *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI)*, 481–88. https://arxiv.org/abs/1301.3886.

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.

Sun, Wei, Robin Hanson, Kathryn B. Laskey, and Charles Twardy. 2012. "Probability and Asset Updating Using Bayesian Networks for Combinatorial Prediction Markets." *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, 815–24. https://arxiv.org/abs/1210.4900.

Xia, Lirong, and David M. Pennock. 2011. "Price Updating in Combinatorial Prediction Markets with Bayesian Networks." *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI)*, 825–34. https://arxiv.org/abs/1202.3756.