

Uniform-Loss Automated Market Making for Prediction Markets

Ciamac C. Moallemi

Decision, Risk, and Operations Division, Graduate School of Business, Columbia University,
New York, NY, USA

Paradigm, New York, NY, USA

Uniswap Labs, New York, NY, USA

Dan Robinson

Paradigm, New York, NY, USA

Brian Zhu

Department of Industrial Engineering and Operations Research, Columbia University,
New York, NY, USA

Abstract

Automated market makers (AMMs) for prediction markets descend from market scoring rules, where a mechanism operator subsidizes a market to aggregate beliefs about uncertain events. The existing literature has focused on bounding the total worst-case loss to the subsidizer, but has not addressed how that loss is distributed across price states or over time. We use the framework of loss-versus-rebalancing (LVR) to study this distribution and introduce *uniform AMMs*, defined by the property that instantaneous LVR is proportional to pool value and independent of the current token price. In a static setting, we show that for a broad class of *win-martingales* – processes that converge to 0 or 1 at a fixed resolution time – there exists a pricing function that achieves uniform LVR under that process, and conversely, that any sufficiently regular pricing function induces a win-martingale under which it is uniform. We then extend the framework to dynamic liquidity management, showing that liquidity levels can be adjusted over time to implement a prescribed target expected cumulative loss schedule. This theory is illustrated with canonical examples of win-martingales and pricing functions. Our results can inform AMM designers and liquidity providers on how the inevitable cost of subsidizing price discovery can be shaped and controlled across both price and time.

2012 ACM Subject Classification Applied computing → Economics

Keywords and phrases Prediction markets, automated market makers, loss-versus-rebalancing, win-martingales, liquidity provision

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.14

Category Extended Abstract

Related Version *Full Version*: <https://arxiv.org/abs/2607.17428>

Funding *Ciamac C. Moallemi*: CCM is supported by the Briger Family Digital Finance Lab at Columbia University. CCM is a research advisor for Paradigm and Uniswap Labs.

Binary prediction markets trade YES and NO claims whose prices sum to one and which resolve to (1,0) or (0,1). Market scoring rules and their equivalent invariant-based AMMs make these markets liquid by committing capital that informed traders can earn as they move stale quotes toward fair value. Classical bounded-loss guarantees control *how much* this subsidy can cost, but not *where* in the price space or *when* before resolution it is spent.

LVR measures the predictable transfer from liquidity providers (LPs) in an AMM to arbitrageurs Milionis et al. [3], and provides the framework to quantify where and when losses occur, and can also be interpreted as an LP’s subsidy for information discovery. We



© Ciamac C. Moallemi, Dan Robinson, and Brian Zhu;
licensed under Creative Commons License CC-BY 4.0
8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyropoulou; Article No. 14; pp. 14:1–14:3

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

propose *uniform LVR* as a desiderata of AMM design for prediction markets (or binary option markets in general): after accounting for the common rate of information arrival, every probability state incurs the same instantaneous LVR per dollar of pool value. Besides resulting in state-independent losses and thus removing price-based timing incentives for LPs, we show that uniform LVR minimizes a worst-case relative loss criterion across price states.

We assume that the fair YES price $P_t \in [0, 1]$ is a continuous *win-martingale*, i.e., P is a martingale and $P_T \in \{0, 1\}$ almost surely at resolution time T . To balance breadth with tractability, we study a class of separable dynamics:

$$dP_t = \frac{G(P_t)}{h(t)} dW_t, \quad \int_0^t h(s)^{-2} ds < \infty \ (t < T), \quad \int_0^T h(s)^{-2} ds = \infty. \quad (1)$$

Here G determines where beliefs are most responsive to information, while h determines when information arrives. The divergent information clock guarantees binary resolution under standard regularity and absorbing-boundary conditions on G .

We represent an AMM by the value $V(p)$ of its arbitrated reserves at prices $(p, 1-p)$. The AMM's YES and NO reserves are $X(p) = V(p) + (1-p) \cdot V'(p)$ and $Y(p) = V(p) - p \cdot V'(p)$, respectively. Itô's formula yields pool-value dynamics and instantaneous LVR rate:

$$dV(P_t) = V'(P_t) \cdot \frac{G(P_t)}{h(t)} dW_t - \text{LVR}_t dt, \quad \text{LVR}_t = -\frac{1}{2} \cdot V''(P_t) \cdot \frac{G(P_t)^2}{h(t)^2}. \quad (2)$$

Uniformity requires $\text{LVR}_t/V(P_t)$, after normalizing by the information clock $h(t)^{-2}$, to be constant in P_t . Equivalently, the pool-value function solves the boundary value problem

$$G(p)^2 \cdot V''(p) + \beta \cdot V(p) = 0, \quad V(0) = V(1) = 0, \quad V(p) > 0 \ (0 < p < 1). \quad (3)$$

Under mild endpoint regularity, this singular Sturm–Liouville problem has a positive, concave, and unimodal solution. Conversely, any sufficiently regular strictly concave V induces a win-martingale under which V has uniform LVR through the volatility profile


$$G(p) = \sqrt{\beta \cdot \frac{V(p)}{-V''(p)}}. \quad (4)$$

This is the paper's central bidirectional correspondence: underlying belief dynamics determine the AMM geometry that equalizes losses, while an existing invariant reveals the belief process under which it is uniform. Canonical examples illustrate this correspondence; for example:


- Logistic volatility $G(p) = p(1-p)$ yields $V(p) = 2\sqrt{p(1-p)}$, i.e., the constant-product invariant $xy = L^2$.
- Wright–Fisher volatility $G(p) = \sqrt{p(1-p)}$ yields $V(p) = p(1-p)$, i.e., the constant-elasticity invariant $\sqrt{x} + \sqrt{y} = \sqrt{L}$.
- Gaussian score dynamics Archak and Ipeiritos [1], $G(p) = \phi(\Phi^{-1}(p))$ yield $V(p) = \phi(\Phi^{-1}(p))$, the pm-AMM in [4].
- The “most exciting game” Backhoff-Veraguas and Beiglböck [2] has volatility profile $G(p) = \sin(\pi p)/\pi$ and yields a family of solutions for V based on Ferrers functions.
- In the converse direction, the LMSR value $V(p) = L[-p \log p - (1-p) \log(1-p)]$ is uniform for $G(p) = \sqrt{p(1-p) \cdot H(p)}$, where H is the binary entropy function.

Finally, the invariant controls the subsidy across *states*, while a deterministic liquidity scale L_t controls it across *time*. For any admissible differentiable, nondecreasing target expected cumulative loss fraction $D(t)$ and initial wealth w_0 , the schedule

$$L_t = \frac{2w_0}{\beta V(P_0)} h(t)^2 \exp\left(\int_0^t \frac{\beta}{2h(s)^2} ds\right) D'(t) \quad (5)$$



■ **Figure 1** Normalized pool-value functions (left) and their reserve invariants (right). The same value profile can be translated into reserve demands.



■ **Figure 2** Left: YES-side market depth $|X'(p)|$, the reserve quantity available per unit change in the implied YES probability. Right: Uniswap-v3-equivalent *liquidity fingerprints*. For relative price $q = p/(1 - p)$, the fingerprint translates local reserve response into the active liquidity $\ell(q) = -2q^{3/2}dX/dq$ that would reproduce the same infinitesimal price impact in Uniswap v3, normalized by the constant-product benchmark. It exposes where an invariant concentrates effective liquidity even when it is not implemented in Uniswap v3.

implements $\mathbb{E}[W_t] = w_0(1 - D(t))$ while preserving $\text{LVR}_t = \beta L_t V(P_t)/[2h(t)^2]$, subject to initial cash feasibility. Hence the same uniform invariant can support front-loaded or back-loaded price-discovery programs without changing which probability states are subsidized. Read the full paper at <https://arxiv.org/abs/2607.17428>.

References

- 1 Nikolay Archak and Panagiotis G. Ipeiritis. Modeling volatility in prediction markets. In *Proceedings of the 10th ACM Conference on Electronic Commerce, EC '09*, pages 231–238, New York, NY, USA, 2009. Association for Computing Machinery. doi:10.1145/1566374.1566414.
- 2 Julio Backhoff-Veraguas and Mathias Beiglböck. The most exciting game. *arXiv preprint arXiv:2305.14037*, 2023. arXiv:2305.14037.
- 3 Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated market making and loss-versus-rebalancing. *arXiv preprint arXiv:2208.06046*, 2022. arXiv:2208.06046.
- 4 Ciamac Moallemi and Dan Robinson. pm-amm: A uniform amm for prediction markets. *Paradigm Research Blog Post*, 2024.