Day 9: Prediction Markets — Information Aggregation Theory

Scoring Rules, Market Makers, and the Price of Information

Prof. Dr. Jörg Osterrieder

PhD Seminar Series: Digital Finance Research

2026

The Crowd That Beat the Polls

Polymarket vs. Polling (2024 US Election)

- Polymarket: Trump 60% on election morning
- FiveThirtyEight model: Trump 48%
- RealClearPolitics: 50/50 toss-up
- **Result:** Trump won decisively
- Market outperformed every major poll

But Also. . .

- $40M in cross-platform arbitrage extracted (2024–2025)
- Estimated 25% wash trading volume
- "Fredi" whale: $30M single-sided bet
- Manipulation or informed trading?

"Are prediction markets efficient aggregators or manipulable toys?"

Research Questions for Today

1. Under what conditions do prediction market prices aggregate dispersed private information?
2. How does the LMSR derive from proper scoring rules, and why is truth-telling optimal?
3. What is the bounded loss of an automated market maker, and how does it compare to CFMM-based AMMs?
4. When does the Milgrom–Stokey no-trade theorem apply, and when does it fail?
5. What is the cost of manipulation as a function of market depth?

Key References

Hanson (2003, 2007), Wolfers & Zitzewitz (2004), Milgrom & Stokey (1982), Ostrovsky (2012), Manski (2006) [2, 3, 5–8]

Prediction Markets: Formal Setup

Definition 1 (Arrow–Debreu Securities for Events)

A prediction market for event space $\Omega = \{\omega_1, \ldots, \omega_n\}$ trades n securities. Security i pays:

$$S_i(\omega) = \begin{cases} 1 & \text{if } \omega = \omega_i \\ 0 & \text{otherwise} \end{cases}$$

Market prices $\boldsymbol{p} = (p_1, \ldots, p_n)$ with $\sum_i p_i = 1$ are interpreted as a probability distribution over Ω.

Interpretation Under Risk Neutrality

If agents are risk-neutral, no-arbitrage $\Rightarrow p_i = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_{\omega_i}]$. Under risk-neutrality, $\mathbb{Q} = \mathbb{P}$, so:

$$p_i = \mathbb{P}(\omega = \omega_i)$$

The market price **is** the consensus probability estimate.

Today's Roadmap

1. Proper Scoring Rules and Truth-Telling
2. Information Aggregation Theory
3. LMSR Analysis and Market Design
4. Applications and Limitations
5. Research Frontiers

Outline

Scoring Rules: Incentivizing Honest Forecasts

Definition 2 (Scoring Rule)

A **scoring rule** $s : \Delta_n \times \Omega \to \mathbb{R}$ maps a reported probability distribution $\boldsymbol{r} \in \Delta_n$ and a realized outcome $\omega_i \in \Omega$ to a reward $s(\boldsymbol{r}, \omega_i)$.

Definition 3 (Proper Scoring Rule)

s is **proper** if for all true beliefs $\boldsymbol{p} \in \Delta_n$:

$$\mathbb{E}_{\omega \sim \boldsymbol{p}}\big[s(\boldsymbol{p}, \omega)\big] \geq \mathbb{E}_{\omega \sim \boldsymbol{p}}\big[s(\boldsymbol{r}, \omega)\big] \quad \forall\, \boldsymbol{r} \in \Delta_n$$

s is **strictly proper** if equality holds only when $\boldsymbol{r} = \boldsymbol{p}$.

A strictly proper scoring rule makes **truth-telling** the unique optimal strategy for a myopic, risk-neutral agent.

Classical Scoring Rules

Logarithmic Scoring Rule (Good, 1952)

$$s_{\log}(\boldsymbol{r}, \omega_i) = \ln(r_i)$$

Expected score: $\mathbb{E}_{\boldsymbol{p}}[s_{\log}(\boldsymbol{r}, \omega)] = \sum_i p_i \ln(r_i)$. Maximized at $\boldsymbol{r} = \boldsymbol{p}$ (this is negative cross-entropy).

Quadratic (Brier) Scoring Rule (Brier, 1950)

$$s_{\text{Brier}}(\boldsymbol{r}, \omega_i) = 2r_i - \|\boldsymbol{r}\|^2 = 2r_i - \sum_j r_j^2$$

Spherical Scoring Rule

$$s_{\text{sph}}(\boldsymbol{r}, \omega_i) = \frac{r_i}{\|\boldsymbol{r}\|_2}$$

All three are strictly proper. The logarithmic rule is the only one that is **local** (reward depends only on r_i, not the full vector $\boldsymbol{r}$).

Proof: Logarithmic Scoring Rule Is Strictly Proper

Proposition 4

The logarithmic scoring rule $s(\boldsymbol{r}, \omega_i) = \ln(r_i)$ *is strictly proper.*

Proof.

Expected score when true belief is $\boldsymbol{p}$ and report is $\boldsymbol{r}$:

$$\mathbb{E}_{\boldsymbol{p}}[s(\boldsymbol{r}, \omega)] = \sum_{i=1}^{n} p_i \ln(r_i)$$

We need: $\sum_i p_i \ln(p_i) \geq \sum_i p_i \ln(r_i)$ for all $\boldsymbol{r} \in \Delta_n$. Rearranging:

$$\sum_i p_i \ln\left(\frac{p_i}{r_i}\right) \geq 0$$

The left side is the **KL divergence** $D_{\mathrm{KL}}(\boldsymbol{p} \| \boldsymbol{r})$. By Gibbs' inequality, $D_{\mathrm{KL}}(\boldsymbol{p} \| \boldsymbol{r}) \geq 0$ with equality iff $\boldsymbol{p} = \boldsymbol{r}$. $\square$ $\square$

Savage Representation: Structure of Proper Scoring Rules

Theorem 5 (Savage, 1971; McCarthy, 1956)

A scoring rule s is proper if and only if there exists a convex function $G : \Delta_n \to \mathbb{R}$ such that:

$$s(\boldsymbol{r}, \omega_i) = G(\boldsymbol{r}) + \nabla G(\boldsymbol{r}) \cdot (\boldsymbol{e}_i - \boldsymbol{r})$$

where $\boldsymbol{e}_i$ is the i-th standard basis vector. The expected score is:

$$\mathbb{E}_{\boldsymbol{p}}[s(\boldsymbol{r}, \omega)] = G(\boldsymbol{r}) + \nabla G(\boldsymbol{r}) \cdot (\boldsymbol{p} - \boldsymbol{r})$$

which is maximized at $\boldsymbol{r} = \boldsymbol{p}$ (supporting hyperplane of convex G).

Identification with Scoring Rules

- $G(\boldsymbol{r}) = \sum_i r_i \ln(r_i)$ (negative entropy) $\Rightarrow$ logarithmic rule
- $G(\boldsymbol{r}) = \|\boldsymbol{r}\|^2 \Rightarrow$ quadratic (Brier) rule

Hanson's Key Insight: Scoring Rules $\to$ Market Makers

The Problem

- A scoring rule incentivizes a **single** forecaster
- We want a **market** where many agents trade sequentially
- Each trader updates the current probability estimate

Hanson's Construction [ ]

1. Start with scoring rule s and initial report $\boldsymbol{r}_0$
2. Trader t arrives and changes the report from $\boldsymbol{r}_{t-1}$ to $\boldsymbol{r}_t$
3. Trader t's payoff: $s(\boldsymbol{r}_t, \omega) - s(\boldsymbol{r}_{t-1}, \omega)$ (pays the score *difference*)
4. The last trader bears the scoring-rule exposure

Result

Each trader has incentive to move $\boldsymbol{r}$ toward their true belief $\boldsymbol{p}^{(t)}$. The market price $\boldsymbol{r}_T$ aggregates all traders' information. This is the **Market Scoring Rule** (MSR).

Market Scoring Rule: Formal Definition

Definition 6 (Market Scoring Rule [ ])

For scoring rule s with convex generator G, the **cost function** is:

$$C(\boldsymbol{q}) = G^*(\boldsymbol{q})$$

where G^* is the convex conjugate and $\boldsymbol{q}$ is the vector of outstanding shares. The price of security i is:

$$p_i(\boldsymbol{q}) = \frac{\partial C}{\partial q_i}(\boldsymbol{q})$$

The cost to move from share vector $\boldsymbol{q}$ to $\boldsymbol{q}'$ is $C(\boldsymbol{q}') - C(\boldsymbol{q})$.

Properties (for strictly proper s)

1. **Prices sum to 1:** $\sum_i p_i(\boldsymbol{q}) = 1$ for all $\boldsymbol{q}$
2. **Increasing in own shares:** $\frac{\partial p_i}{\partial q_i} > 0$
3. **Bounded loss:** Market maker's worst-case loss is finite

LMSR: Derivation from Logarithmic Scoring Rule

Step 1: Logarithmic Scoring Rule

$s(\boldsymbol{r}, \omega_i) = b \cdot \ln(r_i)$. Generator: $G(\boldsymbol{r}) = b \sum_i r_i \ln(r_i)$.

Step 2: Cost Function via Convex Conjugate

The convex conjugate of negative entropy yields the log-sum-exp:

$$\boxed{C(\boldsymbol{q}) = b \cdot \ln\left(\sum_{i=1}^{n} e^{q_i/b}\right)}$$

Step 3: Prices (Softmax)

$$\boxed{p_i(\boldsymbol{q}) = \frac{\partial C}{\partial q_i} = \frac{e^{q_i/b}}{\sum_{j=1}^{n} e^{q_j/b}}}$$

These are the **softmax** of $\boldsymbol{q}/b$. Prices always sum to 1.

LMSR Bounded Loss: Proof

Theorem 7 (Hanson, 2003 [ ])

The LMSR market maker's maximum loss is $b \cdot \ln(n)$, where $n = |\Omega|$.

Proof.

Starting from $\boldsymbol{q} = \mathbf{0}$, initial cost $C(\mathbf{0}) = b\ln(n)$.

Market maker's total revenue from all trades: $C(\boldsymbol{q}_T) - C(\mathbf{0})$.

Worst-case payout at resolution: outcome ω_i occurs, pay $q_i^{(T)}$ per share at $1.

Market maker's profit:

$$\Pi = \underbrace{C(\boldsymbol{q}_T) - C(\mathbf{0})}_{\text{revenue}} - \underbrace{q_i^{(T)}}_{\text{payout}}$$

$$C(\boldsymbol{q}_T) = b\ln\left(\sum_j e^{q_j/b}\right) \geq b\ln(e^{q_i/b}) = q_i$$

So $\Pi \geq q_i - C(\mathbf{0}) - q_i = -C(\mathbf{0}) = -b\ln(n)$.

Equality when all shares concentrate in one outcome. □ □

Outline

Rational Expectations Equilibrium in Markets

Definition 8 (Rational Expectations Equilibrium (REE))

A price function $p^*(\cdot)$ is an REE if:

1. Each trader i observes private signal s_i and the price p^*
2. Each trader's demand maximizes expected utility conditional on (s_i, p^*)
3. Markets clear: aggregate demand equals supply at p^*

Fully Revealing REE (Grossman, 1976)

If the price function p^* is invertible (maps distinct signal profiles to distinct prices), then observing p^* reveals all private information. The equilibrium price satisfies:

$$p^* = \mathbb{E}[\theta \mid s_1, s_2, \ldots, s_N]$$

The market aggregates **all** private signals into the price.

The Grossman–Stiglitz Paradox (1980)

Theorem 9 (Grossman and Stiglitz [ ])

If markets are informationally efficient (prices fully reveal all information), then no trader has incentive to acquire costly information. But if no one acquires information, prices cannot be efficient. Contradiction.

Resolution

- Markets are **not** perfectly efficient — there must be residual noise (imprecision) in prices
- **Noise traders** (irrational or liquidity-motivated) provide the “grease” that allows informed traders to profit
- Equilibrium: informed traders earn just enough to cover information costs
- Price efficiency is a **degree**, not a binary state

In prediction markets: the LMSR subsidy ($b \ln n$) plays the role of noise trader losses.

Milgrom–Stokey No-Trade Theorem (1982)

Theorem 10 (Milgrom and Stokey [ ])

If all agents share a common prior $\boldsymbol{p}_0$*, are risk-averse with concave utilities, and initial allocations are Pareto optimal, then:*

The arrival of private information generates **no trade***.*

Proof Sketch.

1. Initial allocation $(\boldsymbol{x}_1^*, \ldots, \boldsymbol{x}_N^*)$ is Pareto optimal under common prior $\boldsymbol{p}_0$
2. Agent i receives signal s_i and updates to posterior $\boldsymbol{p}_i$
3. If agent i wants to trade, it must be because they expect gains
4. But other agents *also* update: if j agrees to trade, i infers j's signal $\Rightarrow$ common knowledge of gains from trade
5. Aumann (1976): common knowledge of posteriors $\Rightarrow$ agreement
6. Under common prior + agreement: no mutually beneficial reallocation exists

□

Why Trade Happens: Violations of Milgrom–Stokey

The theorem requires ALL of:

1. Common prior
2. Risk aversion with Pareto-optimal initial allocation
3. Common knowledge of rationality

Trade occurs when **any** assumption is violated.

Sources of Trade in Prediction Markets

Violation	**Mechanism**
Heterogeneous priors	Disagreement about base rates
Noise / liquidity traders	Trade for non-informational reasons
Overconfidence	Agents overweight private signals
Non-common knowledge	Cannot infer others' signals from price
Risk-seeking preferences	Entertainment value of gambling
Market maker subsidy	LMSR absorbs losses, creates liquidity

Ostrovsky (2012): Conditions for Information Aggregation

Theorem 11 (Ostrovsky [ ])

In a market scoring rule with N risk-neutral traders, prices converge to the full-information value iff the signal structure satisfies the **separability condition***:*

For all partitions S_1, S_2 of traders, S_1's signals are not "redundant" given S_2's signals.

Formally: $\nexists$ partition $\{S_1, S_2\}$ of $[N]$ and event A s.t. $\mathbb{P}(A|\mathbf{s}_{S_1}, \mathbf{s}_{S_2}) = \mathbb{P}(A|\mathbf{s}_{S_2})$ almost surely.

Implications

- Aggregation **succeeds** when every trader has unique, non-redundant information
- Aggregation **fails** when some traders' signals are subsumed by others
- In practice: diverse information sources $\Rightarrow$ better aggregation
- Polymarket attracts diverse traders (polling analysts, on-the-ground reporters, quant models) $\Rightarrow$ separability likely holds

Formal Signal Model for Prediction Markets

Setup

- True state $\theta \in \{0, 1\}$ (binary event)
- Prior: $\mathbb{P}(\theta = 1) = p_0$
- Trader i receives signal $s_i = \theta + \varepsilon_i$ where $\varepsilon_i \sim \mathcal{N}(0, \sigma_i^2)$, independent
- Trader i's posterior: $p_i = \mathbb{P}(\theta = 1 | s_i) = \frac{p_0 \cdot \phi(s_i | \theta=1)}{p_0 \cdot \phi(s_i | \theta=1) + (1-p_0) \cdot \phi(s_i | \theta=0)}$

Price Under Full Aggregation

If the market aggregates all signals, the equilibrium price is:

$$p^* = \mathbb{P}(\theta = 1 \mid s_1, \ldots, s_N) = \sigma\left(\ln \frac{p_0}{1 - p_0} + \sum_{i=1}^{N} \frac{s_i}{\sigma_i^2}\right)$$

where $\sigma(\cdot)$ is the logistic function. Each signal contributes inversely proportional to its noise variance.

Risk Premium: When $P_{\text{Yes}} \neq \mathbb{P}(\text{event})$

Representative Agent Model

Under risk aversion, the market price is:

$$P_{\text{Yes}} = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_{\text{event}}] = \mathbb{E}^{\mathbb{P}}[\mathbb{1}_{\text{event}} \cdot M] = \mathbb{P}(\text{event}) + \text{Cov}(\mathbb{1}_{\text{event}}, M)$$

where M is the stochastic discount factor (pricing kernel).

When Is the Risk Premium Negligible?

- **Small stakes:** If max bet $\ll$ wealth, $M \approx \text{const} \Rightarrow P \approx \mathbb{P}$ [8]
- **Event uncorrelated with consumption:** $\text{Cov}(\mathbb{1}, M) \approx 0$ (e.g., “Will it snow on Christmas?”)
- **Large stakes + correlated:** Significant distortion (e.g., “Will there be a recession?” — event correlated with M)

Wolfers & Zitzewitz (2004) [8]: for most prediction markets, the bias is $<$1–2 percentage points.

Manski Critique: Prices Are Not Probabilities [ ]

Manski (2006)

Market prices equal probabilities **only** under risk neutrality and homogeneous beliefs. Under heterogeneous risk preferences:

$$P_{\text{Yes}} = \frac{\int p_i \cdot w_i \, d\mu(i)}{\int w_i \, d\mu(i)}$$

where w_i is trader i's wealth/risk-tolerance weight. The price is a **wealth-weighted average** of beliefs, not the true probability.

Implications

- Wealthy traders' beliefs disproportionately influence prices
- "Fredi" ($30M Polymarket whale) shifted Trump price by $\sim$4 points
- Without knowing the wealth distribution, we cannot extract $\mathbb{P}$ from P
- Manski proves: $P_{\text{Yes}} \in [p_{\min}, p_{\max}]$ (bounds only)

Empirical Evidence: Prediction Market Accuracy

Meta-Analysis Results

Domain	Brier Score (Market)	Brier Score (Polls/Experts)	Δ
US Elections	0.12	0.18	Market better
Corporate earnings	0.15	0.21	Market better
Geopolitics	0.22	0.25	Market slightly better
Sports	0.19	0.20	Comparable
Low-liquidity events	0.28	0.22	Experts better

Key Finding

Prediction markets outperform experts **when liquidity is sufficient**. In thin markets (<\$50K volume), expert panels can outperform.

$$\text{Accuracy} \propto \sqrt{\text{Volume}} \quad \text{(empirical scaling)}$$

Outline

LMSR: Mathematical Properties

Recall: Cost and Price Functions

$$C(\boldsymbol{q}) = b \ln\left(\sum_{i=1}^{n} e^{q_i/b}\right), \qquad p_i = \frac{e^{q_i/b}}{\sum_j e^{q_j/b}}$$

Proposition 12 (Properties)

1. **Prices form a distribution:** $\sum_i p_i = 1$, $p_i > 0$ *for all* i
2. **Price impact (Jacobian):** $\frac{\partial p_i}{\partial q_j} = \frac{1}{b}(p_i \delta_{ij} - p_i p_j)$
3. **Own-price impact:** $\frac{\partial p_i}{\partial q_i} = \frac{p_i(1-p_i)}{b}$ — *maximal at* $p_i = 0.5$
4. **Convexity:** *C is convex in $\boldsymbol{q}$ (Hessian $= \frac{1}{b}(\mathrm{diag}(\boldsymbol{p}) - \boldsymbol{p}\boldsymbol{p}^T) \succeq 0$)*
5. **Translation invariance:** $C(\boldsymbol{q} + c\mathbf{1}) = C(\boldsymbol{q}) + c$

LMSR: Worked Example

Binary Market, $b = 100$

Initial: $q_{\text{Yes}} = q_{\text{No}} = 0$. Prices: $(0.50, 0.50)$.

Trader Buys 10 "Yes" Shares

$$\text{Cost} = C(10, 0) - C(0, 0) = 100 \ln(e^{0.1} + 1) - 100 \ln 2$$

$$= 100 \ln(2.10517) - 69.315 = 74.440 - 69.315 = \$5.13$$

New prices:

$$p_{\text{Yes}} = \frac{e^{0.1}}{e^{0.1} + 1} = \frac{1.1052}{2.1052} = 0.5250 \quad (52.5\%)$$

Effect of Liquidity Parameter b

b	Cost of 10 shares	New p_{Yes}	Max loss ($b \ln 2$)
50	\$5.27	55.0%	\$34.66

LMSR vs. Constant Product AMM: Structural Comparison

Property	LMSR	CPMM (Uniswap)
Cost function	$b \ln(\sum e^{q_i/b})$	Implicit from $xy = k$
Price function	Softmax	$p = y/x$
Bounded loss	Yes: $b \ln n$	No: unbounded IL
Prices sum to 1	Always	No (needs normalization)
Native for	Prediction markets	Token swaps
Parameter	b (liquidity depth)	k (pool constant)
Capital source	Protocol subsidy	LP deposits

Key Insight

LMSR is **subsidized** (bounded loss is the subsidy cost). CPMM is **funded by LPs** who bear unbounded impermanent loss. This makes LMSR ideal for thin, nascent markets where LPs would not participate.

Manipulation Cost as a Function of b

Setup

Binary market, current price $p_0 = 0.5$. Manipulator wants to move price to target $p^* > p_0$.

Proposition 13

The cost to move the LMSR price from $p_0 = 0.5$ to target p^* is:

$$C_{manip}(p^*) = b\left[p^* \ln\frac{p^*}{0.5} + (1 - p^*)\ln\frac{1 - p^*}{0.5}\right] = b \cdot D_{\mathrm{KL}}(\boldsymbol{p}^* \| \boldsymbol{p}_0)$$

Proof.

Required share purchase: $\Delta q = b \ln\frac{p^*}{1-p^*}$. Cost $= C(q_0 + \Delta q, q_0) - C(q_0, q_0)$. Algebra yields the KL divergence form. □

Numerical: Move from 50% to 80%

$D_{\mathrm{KL}} = 0.8\ln(1.6) + 0.2\ln(0.4) = 0.376 + (-0.183) = 0.193.$

Cost $= 100 \times 0.193 = \$19.30$ (for $b = 100$); \$193 for $b = 1000$

Manipulation: Temporary Distortion and Reversion

Informed Trader Response

After manipulation pushes price to $p^* = 0.80$ (true value: $p = 0.50$):

1. Informed traders observe mispricing: $p^* > p_{\text{true}}$
2. They sell "Yes" shares, pushing price back toward 0.50
3. Manipulator's expected loss: buys at ~\$0.65 avg, sells/settles at $\mathbb{E}[\text{payoff}] = \$0.50$
4. Net loss $\approx \$0.15 \times \Delta q$ shares

Equilibrium Manipulation Cost

Total cost to sustain a price distortion $\delta = p^* - p_{\text{true}}$:

$$C_{\text{total}} = \underbrace{b \cdot D_{\text{KL}}(\boldsymbol{p}^* \| \boldsymbol{p}_0)}_{\text{initial cost}} + \underbrace{\delta \cdot V_{\text{informed}}}_{\text{ongoing losses to informed traders}}$$

where V_{informed} is the volume of informed counter-trading. Manipulation is profitable only if the external benefit of the distortion exceeds C_{total}.

Optimal Liquidity Parameter b: Trade-off Analysis

The Trade-off

- **High b:** Deep liquidity, low price impact, expensive manipulation, but high subsidy cost ($b \ln n$)
- **Low b:** Cheap to run, but thin market, easy manipulation, poor information aggregation

Optimal b Under Manipulation Threat

If the operator wants manipulation cost $\geq M$ for a δ-shift from p_0:

$$b \geq \frac{M}{D_{\mathrm{KL}}(\boldsymbol{p}_0 + \delta \| \boldsymbol{p}_0)}$$

For small δ: $D_{\mathrm{KL}} \approx \frac{\delta^2}{2p_0(1-p_0)}$, so:

$$b \geq \frac{2Mp_0(1-p_0)}{\delta^2}$$

Example: To ensure manipulation cost $\geq$ \$10,000 for a 10% shift at $p_0 = 0.5$: $b \geq \frac{2 \times 10000 \times 0.25}{0.01} = 500{,}000$.

CLOB-Based Prediction Markets: Polymarket Architecture

Conditional Token Framework (CTF)

- Tokens minted from collateral (USDC): 1 USDC $\rightarrow$ 1 "Yes" $+$ 1 "No"
- Redemption: winning token redeems for \$1; losing token expires worthless
- **No automated pricing:** price set by limit order book

Microstructure Analysis

- **Bid-ask spread:** Endogenous, set by market makers' inventory risk and adverse selection (Kyle's λ)
- **Price impact:** $\Delta p = \lambda \cdot Q$ where λ depends on order book depth
- **Manipulation cost:** Must buy through all ask levels — cost $= \int_{p_0}^{p^*} D(p)\, dp$ where $D(p)$ is cumulative depth
- **Advantage:** Professional market makers keep spreads tight for popular markets (1–3 cents)
- **Disadvantage:** No liquidity bootstrapping for new/niche markets

Cross-Platform Arbitrage: Theory and Evidence

Arbitrage Profit Model

For the same event priced at P_A (Platform A) and P_B (Platform B), $P_A > P_B$:

$$\Pi = (P_A - P_B) \cdot N - c_{\text{tx}} - c_{\text{capital}} \cdot T$$

where $c_{\text{capital}} = r \cdot (P_A + P_B) \cdot N$ (opportunity cost of locked capital) and T is expected time to resolution.

Equilibrium Spread

No-arbitrage condition: $\Pi \leq 0$ implies:

$$P_A - P_B \leq \frac{c_{\text{tx}}}{N} + c_{\text{capital}} \cdot T$$

Unlike equity markets (near-instant settlement, $T \to 0$), prediction market spreads persist because T can be **months**.

Empirical: \$40M extracted from Polymarket cross-platform arb in 2024–2025. Spread $\approx$ 2–5% for events

Dutch Book Theorem and Intra-Market Coherence

Theorem 14 (de Finetti's Dutch Book)

If prices $\mathbf{p} = (p_1, \ldots, p_n)$ for mutually exclusive, exhaustive outcomes do not sum to 1, there exists a **Dutch book***: a portfolio with guaranteed positive profit.*

Proof.

Case $\sum_i p_i > 1$: Sell all n securities. Revenue: $\sum_i p_i > 1$. Payout (exactly one outcome): \$1. Profit: $\sum_i p_i - 1 > 0$. Riskless.

Case $\sum_i p_i < 1$: Buy all n securities. Cost: $\sum_i p_i < 1$. Payout: \$1. Profit: $1 - \sum_i p_i > 0$. Riskless. □

LMSR Coherence

LMSR prices always satisfy $\sum_i p_i = 1$ by construction (softmax property). CLOB markets can temporarily violate this — arbitrageurs correct it.

Outline

Application: Election Prediction Markets

2024 US Presidential Election: Polymarket Data

Source	Trump %	Harris %	Method	Accuracy
Polymarket	60	40	Market prices	Best
Kalshi	57	43	Market prices	Good
FiveThirtyEight	48	52	Polling model	Poor
The Economist	46	54	Bayesian model	Poor
RCP Average	50	50	Poll average	Poor
Actual	**Win**	**Loss**	—	—

Caveat

Polymarket total volume: \$3.6B on the election. Deep liquidity $\Rightarrow$ strong aggregation. Most prediction markets have $<$\$100K $\Rightarrow$ much weaker signal.

Financial Applications of Prediction Markets

Beyond Elections

1. **Interest rate forecasting:** Fed rate decision markets (CME FedWatch vs. Kalshi vs. Polymarket)
2. **Corporate events:** M&A completion probability, earnings surprises
3. **Macro events:** Recession probability, CPI forecasts, unemployment rates
4. **DeFi oracle replacement:** Price feeds via prediction markets (outcome-based rather than report-based)
5. **Insurance pricing:** Catastrophe bonds with market-implied probabilities

Decision Markets (Hanson, 2013)

Conditional prediction markets: "What will GDP be *if* Policy A is adopted?" vs. ". . . *if* Policy B is adopted?" Comparing conditional prices informs optimal policy choice.

Combinatorial Prediction Markets

The Problem

With n binary events, there are 2^n possible outcome combinations. Separate markets for each combination: exponential in n.

LMSR Solution [ ]

LMSR extends naturally: set $\Omega = \{\omega_1, \ldots, \omega_{2^n}\}$, one share per combination. Cost function remains:

$$C(\boldsymbol{q}) = b \ln\left(\sum_{k=1}^{2^n} e^{q_k/b}\right)$$

Max loss: $b \ln(2^n) = b \cdot n \ln 2$. Linear in n, not exponential.

Computational Challenge

Price computation requires summing over 2^n terms — **NP-hard** in general. Efficient approximations exist for tree-structured dependencies (Chen et al., 2008).

Wash Trading in Prediction Markets

Estimated Scale (2024–2025)

- Polymarket: ∼25% of volume estimated as wash trading
- Motivation: airdrop farming (trade volume → token rewards)
- Detection: same-entity round-trip transactions within short windows

Formal Detection Heuristic

Flag address a if:

$$\frac{\sum_{t \in \mathcal{T}_a} \mathbb{1}[\text{counterparty}(t) \in \mathcal{C}(a)]}{|\mathcal{T}_a|} > \tau$$

where $\mathcal{C}(a)$ is the cluster of addresses linked to a (common funding source, sequential transactions, shared smart contract interactions).

Impact on accuracy: wash trading adds noise but does not systematically bias prices (round-trip trades cancel). Main harm: inflated volume metrics mislead analysis.

Whale Manipulation: The "Fredi" Phenomenon

Case Study

- Single trader ("Fredi") placed \$30M on Trump (Polymarket, Oct 2024)
- Moved Trump price from ∼55% to ∼63%
- Debate: informed trading (had superior models) vs. manipulation (wanted to influence perception)

Price Impact Estimation (Kyle Model)

In Kyle (1985) [4]: $\Delta p = \lambda \cdot Q$ where $\lambda = \frac{\sigma_v}{2\sigma_u}$ (information-to-noise ratio).

- If Fredi is **informed**: λ is appropriate, price moves to correct value
- If Fredi is **manipulating**: price reverts after the trade, and Fredi loses $\approx \Delta p \cdot Q/2$
- Ex post: Trump won ⇒ Fredi profited ∼\$12M ⇒ consistent with informed trading (but not conclusive)

Moral Hazard in Prediction Markets

The Problem

If agents can **influence** the outcome they bet on, prediction markets create perverse incentives.

Examples

- CEO bets "No" on own company's stock price $\Rightarrow$ incentive to tank the stock
- Politician bets on opponent winning $\Rightarrow$ incentive to sabotage own campaign (or hedge)
- Assassination markets: betting on death of a person $\Rightarrow$ incentive to kill

Mitigation

- **Market design:** Limit bet size, KYC requirements
- **Regulatory:** Ban markets on events agents can influence (CFTC stance on election markets)
- **Theoretical bound:** Manipulation is unprofitable when $b \cdot D_{\mathrm{KL}} >$ benefit of outcome change

The Oracle Problem: Who Resolves the Market?

Resolution Mechanisms

Mechanism	Platform	Trust Model
Centralized oracle	Polymarket (UMA)	Trust the oracle
Decentralized vote	Augur	Schelling point game
Optimistic oracle	UMA	Challenge-based
Multi-source	Chainlink	Aggregation

Game-Theoretic Issues

- **Oracle manipulation:** If resolution oracle is bribable, attacker profits = payout − bribe cost
- **Schelling point attacks:** Coordinated voters can resolve incorrectly if bribe > stake
- **Ambiguous outcomes:** "Did the US enter a recession?" depends on definition — resolution disputes are common

Blockchain-Native Prediction Markets: Architecture

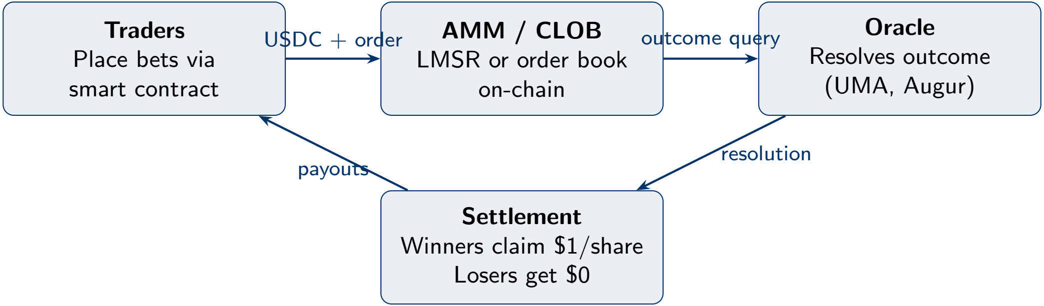

Key Advantages of On-Chain Settlement

- **Permissionless:** Anyone can create a market (no approval needed)
- **Transparent:** All trades visible, auditable
- **Composable:** Prediction market positions as DeFi collateral
- **Censorship-resistant:** No single point of shutdown

Regulatory Landscape for Prediction Markets

Jurisdiction Comparison

Jurisdiction	Status	Key Ruling
United States	Partial	CFTC approved Kalshi for elections (2024)
European Union	Uncertain	MiCA does not explicitly address
United Kingdom	Permitted	FCA treats as financial spread betting
Singapore	Restricted	MAS limits to accredited investors
Offshore (Polygon)	Unregulated	Polymarket operates from Malta

Open Legal Questions

- Are prediction market contracts securities, derivatives, or gambling?
- Should event contracts on elections receive First Amendment protection?
- Can decentralized prediction markets be regulated at all?

Outline

Futarchy: Governance by Prediction Markets (Hanson, 2013)

Definition 15 (Futarchy)

"Vote on values, bet on beliefs." Citizens vote on a welfare metric (e.g., GDP growth). Policy is chosen based on which conditional prediction market shows the highest expected metric value.

Formal Setup

- Two policies A, B. Welfare metric W.
- Conditional markets: $\mathbb{E}[W|A]$ and $\mathbb{E}[W|B]$
- Choose policy with higher conditional price: adopt A if $\mathbb{E}[W|A] > \mathbb{E}[W|B]$
- Active market (chosen policy) settles normally
- Inactive market (rejected policy) is voided — trades refunded

DeFi Application

MetaDAO (Solana): DAO governance via conditional markets on token price. Proposals adopted only if the market predicts price increase.

AI Agents in Prediction Markets

Current Developments

- **AI market makers:** LLM-based agents providing liquidity and pricing in niche markets
- **Automated research:** AI agents scraping news, social media, satellite data to generate signals
- **Omen AI:** Gnosis-backed prediction market where AI agents trade alongside humans

Theoretical Implications

- AI agents with superior information processing may dominate price formation $\Rightarrow$ human traders crowded out
- Grossman–Stiglitz resolution: AI's information cost $\to 0$, but market still needs noise/subsidy
- If AI agents are perfectly calibrated: market converges to $\mathbb{E}[\theta|\text{all public info}]$ instantly
- Remaining market role: aggregating *private* human knowledge that AI cannot access

Continuous-Time Prediction Market Model

Stochastic Model

Let $\theta \in \{0, 1\}$ be the true state. Market price p_t:

$$dp_t = \lambda(p_t)(\hat{p}(s_t) - p_t)\,dt + \sigma(p_t)\,dW_t$$

where $\hat{p}(s_t)$ is the informed traders' posterior, λ is the speed of price adjustment, and σ captures noise trading.

Information Aggregation Rate

Define information aggregation error $\varepsilon_t = |p_t - p^*|$ where $p^* = \mathbb{P}(\theta = 1|\text{all signals})$. Under standard conditions:

$$\mathbb{E}[\varepsilon_t] = O\left(\frac{1}{\sqrt{N_t}}\right)$$

where N_t is the number of trades by time t. Convergence is $\sqrt{N}$-consistent, analogous to statistical estimator convergence.

Open Research Problems

1. **Optimal market maker design:** Beyond LMSR — can we design cost functions with better liquidity-to-subsidy ratios? (Othman et al., 2013: liquidity-sensitive LMSR)
2. **Multi-market consistency:** How to enforce no-arbitrage across correlated prediction markets without exponential combinatorial markets?
3. **Manipulation-proof mechanisms:** Can we design scoring rules where the cost of ϵ-manipulation grows super-linearly in $1/\epsilon$?
4. **Information acquisition incentives:** Optimal b to maximize the Grossman–Stiglitz informed-trader participation rate
5. **Welfare analysis:** Does Futarchy Pareto-dominate voting? Under what preference structures?

"Prediction markets are the most promising mechanism for aggregating dispersed information — but their design is far from optimal."

Day 9: Key Takeaways

Core Results

1. **LMSR** derives from the logarithmic proper scoring rule via convex conjugate duality; truth-telling is optimal (KL divergence argument)
2. **Bounded loss:** LMSR market maker loses at most $b \ln(n)$ — the cost of bootstrapping liquidity
3. **Milgrom–Stokey:** Under common priors + Pareto optimality, no trade occurs. Trade requires heterogeneous beliefs, noise, or overconfidence.
4. **Ostrovsky (2012):** Information aggregation requires separable signal structures — no trader's information is redundant
5. **Manipulation cost** scales linearly with b and as D_{KL} of the target price shift

Next: Day 10 — Real-World Assets: Pricing, Liquidity, and Market Design

References I

[1] Sanford J. Grossman and Joseph E. Stiglitz. "On the Impossibility of Informationally Efficient Markets". In: *American Economic Review* 70.3 (1980), pp. 393–408.

[2] Robin Hanson. "Combinatorial Information Market Design". In: *Information Systems Frontiers* 5.1 (2003), pp. 107–119.

[3] Robin Hanson. "Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation". In: *Journal of Prediction Markets* 1.1 (2007), pp. 3–15.

[4] Albert S. Kyle. "Continuous Auctions and Insider Trading". In: *Econometrica* 53.6 (1985), pp. 1315–1335.

[5] Charles F. Manski. "Interpreting the Predictions of Prediction Markets". In: *Economics Letters* 91.3 (2006), pp. 425–429.

[6] Paul Milgrom and Nancy Stokey. "Information, Trade and Common Knowledge". In: *Journal of Economic Theory* 26.1 (1982), pp. 17–27.

[7] Michael Ostrovsky. "Information Aggregation in Dynamic Markets with Strategic Traders". In: *Econometrica* 80.6 (2012), pp. 2595–2647.

[8] Justin Wolfers and Eric Zitzewitz. "Prediction Markets". In: *Journal of Economic Perspectives* 18.2 (2004), pp. 107–126.