

Self-Financed Wagering Mechanisms for Forecasting

Lambert, Langford, Wortman, Chen, Reeves, Shoham & Pennock

ACM Conference on Electronic Commerce (EC'08)

Outline

- 1 Introduction
- 2 Weighted-Score Mechanisms
- 3 Desirable Properties
- 4 Uniqueness
- 5 Further Extensions

Motivating Scenarios — General Wagering Mechanisms

- March madness
- Fantasy football
- Betting on the weather tomorrow
- Polymarket
- Card games
- Roulette, craps, various casino games
- Horse races

The Formal Model: Wagering Mechanisms

The Setting

- Finite set of **players** $\mathcal{N} = \{1, \dots, n\}$.
- A **random experiment** with outcome $\omega \in \Omega$.
- Each player i holds a private **belief** P_i (a probability distribution over Ω).

Definition 1: Wagering Mechanism

A **wagering mechanism** is a tuple $(\mathcal{R}, \Omega, \Pi)$ with players \mathcal{N} , where:

- \mathcal{R} is the set of allowed **reports**.
- Ω is the set of **outcomes**.
- $\Pi = (\Pi_i)_{i \in \mathcal{N}}$ are **payout functions** $\Pi_i : \mathcal{R}^{\mathcal{N}} \times \mathbb{R}_+^{\mathcal{N}} \times \Omega \mapsto \mathbb{R}_+$.
- $\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = 0$ whenever $m_i = 0$.

Player i 's **net profit** is $\Pi_i(\mathbf{r}, \mathbf{m}, \omega) - m_i$. We denote the sum of the wagers $M = \sum_j m_j$.

Scoring Function Recap: Why We Start Here

What is a scoring function?

A **scoring function** $s(r, \omega)$ assigns a numerical reward based on:

- the player's **report** r , and
- the realized **outcome** ω .

The better the report matches the realized outcome, the higher the score.

Strictly Proper Scoring Rule

A score function s is **strictly proper** for the target property Γ if, for every belief P ,

$$\Gamma(P) = \operatorname{argmax}_{r \in \mathcal{R}} \mathbb{E}_{\omega \sim P}[s(r, \omega)].$$

Why This Alone Is Not Enough

Proper scoring rules give **truthful incentives**, but usually require an **external patron** to fund the elicitation. This paper asks: can we keep truthfulness *without* needing outside money?

The Weighted-Score Mechanism: Definition

Definition 3: Weighted-Score Mechanism (WSM)

A **weighted-score mechanism (WSM)** uses a strictly proper score function $s : \mathcal{R} \times \Omega \rightarrow [0, 1]$, with payout:

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = m_i \left(1 + s(r_i, \omega) - \frac{1}{M} \sum_{j=1}^n m_j s(r_j, \omega) \right)$$

Previous works, and new grounds covered:

Kilgour & Gerchak (2004) adapt scoring rules to be budget-balanced but require *equal wagers* from all players. **Johnstone (2007)** builds off the above; they allow different wagers but produces a mechanism that is not truthful. This paper addresses both simultaneously using this special mechanism.

The Weighted-Score Mechanism: Intuition

Reading the payout step by step:

$$\Pi_i = m_i \left(\underbrace{1}_{\text{wager returned}} + \underbrace{s(r_i, \omega)}_{\text{your score}} - \underbrace{\frac{\sum_j m_j s(r_j, \omega)}{\sum_j m_j}}_{\text{wager-weighted avg of all scores}} \right)$$

your "comparative score"

- You always get your wager back as a baseline.
- You are rewarded for outperforming others: if $s(r_i, \omega) > \bar{s}$, the bracket > 1 and you profit.
- Your wager is a scaling factor; increasing m_i increases both potential gains and potential losses.

Example: Sports Pool

Three friends predict $\mathbb{P}(\text{Team Brauniacs wins})$ using the Brier/quadratic score.

	Report r_i	Wager m_i	$s(\text{Brauniacs win})$	$s(\text{Brauny Bunch win})$
Ollie	0.70	\$30	$1 - (0.30)^2 = 0.91$	$1 - (0.70)^2 = 0.51$
Nishant	0.50	\$20	$1 - (0.50)^2 = 0.75$	$1 - (0.50)^2 = 0.75$
Yilun	0.35	\$10	$1 - (0.65)^2 = 0.5775$	$1 - (0.35)^2 = 0.8775$

Assume Brauniacs win. Weighted-average score:

$$\bar{s} = (30 \cdot 0.91 + 20 \cdot 0.75 + 10 \cdot 0.5775) / 60 = 48.075 / 60 \approx 0.8.$$

$$\text{Payouts } \Pi_i = m_i(1 + s(r_i, 1) - \bar{s}) = m_i(s(r_i, 1) + 0.2):$$

$$\Pi_O = 30(0.91 + 0.2) = \$33.3, \quad \Pi_N = 20(0.75 + 0.2) = \$19, \quad \Pi_Y = 10(0.5775 + 0.2) \approx \$7.78.$$

Notice the total: $\$33.3 + \$19 + \$7.78 = \$60.08 \approx \$60$.

Noticing Budget Balance

This raises the questions: what properties could a well-designed wagering mechanism have?
Can we find one that satisfies all of them simultaneously?

What Goals Could We Have for a Wagering Mechanism?

The Design Goal

We want a mechanism that is not only truthful, but also works as a real betting system:

- players may wager more or less depending on their confidence,
- the game should be **self-financed**,
- the system should resist obvious forms of manipulation.

The Paper's Approach

Instead of starting with a formula and hoping it has good behavior, the paper takes an axiomatic approach:

- 1 first list the desired properties,
- 2 ask which mechanisms satisfy each,
- 3 show that the WSM (and only the WSM) satisfies the seven proposed properties.

Seven Desirable Properties: Overview

Property	Intuition
1 Budget Balance	Total payouts = total wagers, always
2 Anonymity	Payouts depend on actions, not on identity
3 Truthfulness	Honest reporting is a dominant strategy
4 Normality	Gaining for yourself means others lose
5 Sybilproofness	Splitting/merging accounts is useless
6 Individual Rationality	You can always guarantee ≥ 0 expected profit
7 Monotonicity	Bigger bets amplify both gains and losses

Notation used throughout: $\Gamma(P)$ = true property value under belief P ; s = strictly proper score for Γ with values in $[0, 1]$; $M = \sum_j m_j$ = total pot; \mathbf{r}_{-i} = all reports except i .

WSM Satisfies All Seven Properties

Theorem 1

All weighted-score mechanisms are:

- Budget-balanced,
- Anonymous,
- Truthful,
- Sybilproof,
- Normal,
- Individually rational,
- and Monotonic.

Property 1: Budget Balance

Definition

A mechanism is **budget-balanced** if for all $\mathbf{r}, \mathbf{m}, \omega$:

$$\sum_{i \in \mathcal{N}} \Pi_i(\mathbf{r}, \mathbf{m}, \omega) = \sum_{i \in \mathcal{N}} m_i.$$

Proof for the Weighted-Score Mechanism

$$\begin{aligned} \sum_i \Pi_i &= \sum_i m_i \left(1 + s(r_i, \omega) - \frac{\sum_j m_j s(r_j, \omega)}{M} \right) \\ &= M + \sum_i m_i s(r_i, \omega) - \frac{M \sum_j m_j s(r_j, \omega)}{M} = M. \quad \square \end{aligned}$$

The two scoring sums cancel exactly, for every $\mathbf{r}, \mathbf{m}, \omega$.

Property 2: Anonymity

Definition

A mechanism is **anonymous** if for any permutation σ of \mathcal{N} , any i , any ω :

$$\Pi_i((r_i)_{i \in \mathcal{N}}, (m_i)_{i \in \mathcal{N}}, \omega) = \Pi_{\sigma(i)}((r_{\sigma^{-1}(j)})_{j \in \mathcal{N}}, (m_{\sigma^{-1}(j)})_{j \in \mathcal{N}}, \omega).$$

Proof for the Weighted-Score Mechanism

$\Pi_i = m_i(1 + s(r_i, \omega) - \sum_j m_j s(r_j, \omega)/M)$ depends on player i only through (r_i, m_i) and the symmetric weighted average. Relabelling all players by σ leaves both (r_i, m_i) and the average unchanged. \square

Property 3: Truthfulness

Definition

A mechanism is **truthful for** Γ if for every player i , belief P , others' reports \mathbf{r}_{-i} , and wagers \mathbf{m} :

$$\mathbb{E}_P[\Pi_i((\mathbf{r}_{-i}, \Gamma(P)), \mathbf{m}, \omega)] > \mathbb{E}_P[\Pi_i((\mathbf{r}_{-i}, r_i), \mathbf{m}, \omega)] \quad \forall r_i \neq \Gamma(P).$$

Reporting the true value is a *dominant strategy*: it beats every other report regardless of what everyone else does.

Proof for the Weighted-Score Mechanism

Taking expectations:

$$\mathbb{E}_P[\Pi_i] = m_i \left(1 + \mathbb{E}_P[s(r_i, \omega)] \left(1 - \frac{m_i}{M} \right) - \frac{\sum_{j \neq i} m_j \mathbb{E}_P[s(r_j, \omega)]}{M} \right).$$

Only $\mathbb{E}_P[s(r_i, \omega)]$ depends on player i 's report, entering with coefficient $(1 - m_i/M) > 0$. Since s is strictly proper, this is uniquely maximised at $r_i = \Gamma(P)$. \square

Property 4: Normality

Definition

A mechanism is **normal** if whenever player i changes their report and their expected payout changes, all other players experience a change in the *opposite direction*.

Intuition. The prediction game is zero-sum competitive. Improving your standing necessarily worsens everyone else's.

Property 4: Normality

Proof for the Weighted-Score Mechanism

Let \tilde{r}_i be a modified report and $\tilde{r} = (r_{-i}, \tilde{r}_i)$ such that

$$\mathbb{E}_P[\Pi_i(\tilde{r}, m, \omega)] > \mathbb{E}_P[\Pi_i(r, m, \omega)].$$

Then

$$\mathbb{E}_P[s(\tilde{r}_i, \omega)] > \mathbb{E}_P[s(r_i, \omega)].$$

For any $j \neq i$,

$$\mathbb{E}_P[\Pi_j(\tilde{r}, m, \omega)] - \mathbb{E}_P[\Pi_j(r, m, \omega)] = -\frac{m_j}{\sum_j m_j} (\mathbb{E}_P[s(\tilde{r}_i, \omega)] - \mathbb{E}_P[s(r_i, \omega)]) < 0.$$

Similarly, if

$$\mathbb{E}_P[\Pi_i(\tilde{r}, m, \omega)] < \mathbb{E}_P[\Pi_i(r, m, \omega)]$$

then

$$\mathbb{E}_P[\Pi_j(\tilde{r}, m, \omega)] - \mathbb{E}_P[\Pi_j(r, m, \omega)] > 0.$$

Property 5: Sybilproofness

Definition

A mechanism is **sybilproof** if for any subset $S \subset \mathcal{N}$ with $r_i = r_j$ for all $i, j \in S$, and any redistribution of wagers within S preserving $\sum_{i \in S} m'_i = \sum_{i \in S} m_i$:

- 1 $\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = \Pi_i(\mathbf{r}, \mathbf{m}', \omega)$ for all $i \notin S$.
- 2 $\sum_{i \in S} \Pi_i(\mathbf{r}, \mathbf{m}, \omega) = \sum_{i \in S} \Pi_i(\mathbf{r}, \mathbf{m}', \omega)$.

Intuition. Alice splits her \$30 into three fake accounts of \$10 each (all reporting $r_A = 0.70$). Sybilproofness: Bob and Carol's payouts are unaffected, and Alice's total payout across all three accounts equals her original single-account payout.

Proof for the Weighted-Score Mechanism

All members of S share report r_S , so $s(r_i, \omega) = s(r_S, \omega)$ for $i \in S$. The weighted average depends on S only through $\sum_{j \in S} m_j$, which is preserved.

Part 1: For $i \notin S$: (r_i, m_i, \bar{s}) all unchanged, so Π_i unchanged. ✓

Part 2: $\sum_{i \in S} \Pi_i = (\sum_{i \in S} m_i)(1 + s(r_S, \omega) - \bar{s})$ depends only on $\sum_{i \in S} m_i$, which is preserved.

□

Property 6: Individual Rationality

Definition

A mechanism is **individually rational** if for every i , belief P , and $m_i > 0$, there exists r_i^* such that for all choices $(\mathbf{m}_{-i}, \mathbf{r}_{-i})$ by others:

$$\mathbb{E}_P[\Pi_i((\mathbf{r}_{-i}, r_i^*), \mathbf{m}, \omega) - m_i] \geq 0.$$

Proof for the Weighted-Score Mechanism

Take $r_i^* = \Gamma(P)$. For any other player j : $\mathbb{E}_P[s(r_j, \omega)] \leq \mathbb{E}_P[s(\Gamma(P), \omega)]$ (strict propriety applied to P). Therefore:

$$\frac{\sum_j m_j \mathbb{E}_P[s(r_j, \omega)]}{M} \leq \mathbb{E}_P[s(\Gamma(P), \omega)].$$

Thus:

$$\mathbb{E}_P[\Pi_i - m_i] = m_i \left(\mathbb{E}_P[s(\Gamma(P), \omega)] - \frac{\sum_j m_j \mathbb{E}_P[s(r_j, \omega)]}{M} \right) \geq 0. \quad \square$$

Property 7: Monotonicity

Definition

A mechanism is **monotonic** if for all i , belief P , reports \mathbf{r} , wagers \mathbf{m} , and any larger wager $M_i > m_i$:

$$\text{if } \mathbb{E}_P[\Pi_i(\mathbf{m}) - m_i] > 0 \Rightarrow \mathbb{E}_P[\Pi_i(M_i) - M_i] > \mathbb{E}_P[\Pi_i(\mathbf{m}) - m_i] > 0,$$

and symmetrically for losses.

Proof for the Weighted-Score Mechanism

Let $\tilde{s}_j = \mathbb{E}_P[s(r_j, \omega)]$ and $b = \sum_{j \neq i} m_j$. Then:

$$\mathbb{E}_P[\Pi_i - m_i] = \frac{a m_i}{m_i + b}, \quad a = \tilde{s}_i b - \sum_{j \neq i} m_j \tilde{s}_j \quad (\text{independent of } m_i).$$

The function $f(x) = ax/(x + b)$ has derivative $ab/(x + b)^2$, same sign as a . So profit is increasing in m_i when $a > 0$ (you beat the field) and decreasing when $a < 0$ (you trail the field). \square

WSM Satisfies All Seven Properties — And Is the Only One That Does

Theorem 3 — The Main Result (Uniqueness)

A wagering mechanism is budget-balanced, anonymous, truthful for Γ , normal, and sybilproof *if and only if* it is a weighted-score mechanism:

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = m_i \left(1 + s(r_i, \omega) - \frac{1}{M} \sum_j m_j s(r_j, \omega) \right),$$

where s is smooth ($s \in C^2$) and strictly proper for Γ with values in $[0, 1]$.

Proof roadmap: impose five properties in bite-sized chunks.



Step 1 (§5.1): Imposing Truthfulness + Normality

Lemma 1

A wagering mechanism is truthful for Γ and normal if and only if its payouts are nonnegative and additively separable in the form

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = m_i + f_{i,i}(r_i, \mathbf{m}, \omega) - \sum_{j \neq i} f_{i,j}(r_j, \mathbf{m}, \omega),$$

where for all i, j , for any fixed value of \mathbf{m} , $f_{i,j} \in C^2$ is a strictly proper scoring function for Γ .

Additively separable \implies **cross partials vanish**. If $\Pi_k(\mathbf{r}, \omega)$ depended jointly on r_i and r_j , i 's optimal report would depend on r_j , breaking dominant-strategy truthfulness.

What changed?

Before: Π_i can have arbitrary interactions between any pair of reports.

After: Π_i decomposes into n separate score functions, one per player. Functions $f_{i,j}$ can still differ across pairs (i, j) .

§5.2: Adding Anonymity

Previous form (truthful + normal):

$$\Pi_i = m_i + f_{i,i}(r_i, \mathbf{m}, \omega) - \sum_{j \neq i} f_{i,j}(r_j, \mathbf{m}, \omega)$$

Lemma 3

If we add anonymity and equal wagers m , total $M = nm$, $\exists f, g \in C^2$ with payout function for player i being:

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = m + f(r_i, m, M, \omega) - \sum_{j \neq i} g(r_j, m, M, \omega),$$

and further, for any fixed m and M , f and g are strictly proper score functions for Γ that *do not* rely on i, j .

Anonymity Intuition and Explanation

What changed, what do f vs. g mean?

Anonymity says the player label is irrelevant. Swapping Austin (i) and Kirill (j) must swap the payout functions too, forcing $f_{i,i} = f_{j,j}$ (call it f) and $f_{i,j} = f_{j,i}$ (call it g) for all pairs. All potentially n^2 distinct functions collapse to just **two**:

- f : how much your good performance benefits you.
- g : how much another player's good performance hurts you.

These two functions are not yet forced to be equal. Budget balance will pin down their exact amounts.

§5.3: Adding Budget Balance

Previous form (truthful + normal + anonymous + equal wagers):

$$\Pi_i = m + f(r_i, m, M, \omega) - \sum_{j \neq i} g(r_j, m, M, \omega)$$

Lemma 4

If we add budget-balanced, $\exists f \in C^2$ such that payouts are

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = m + f(r_i, m, M, \omega) - \frac{1}{n-1} \sum_{j \neq i} f(r_j, m, M, \omega),$$

furthermore, for any fixed values of m and M , f must be a strictly proper score function for Γ .

§5.4: Sybilproofness

Where we are. After budget balance + anonymity + truthfulness + normality + equal wagers:

$$\Pi_i = m + f(r_i, m, M, \omega) - \frac{1}{n-1} \sum_{j \neq i} f(r_j, m, M, \omega).$$

Real mechanisms must handle *arbitrary* wagers.

The Role of Sybilproofness

It says: a player wagering m_i is indistinguishable from k clones each wagering m_i/k with the same report. This means that we can now allow for arbitrary wagers, and split the wager for each player down into many clones until we obtain equal wagers, allowing us to use that handy proof technique of equal wagers.

Main Uniqueness Result

Theorem 3

A wagering mechanism is budget-balanced, anonymous, truthful, normal, and sybilproof *if and only if* the payouts are given by

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = m_i \left(1 + s^M(r_i, \omega) - \frac{1}{M} \sum_j m_j s^M(r_j, \omega) \right)$$

where s^M is a smooth function taking values in $[0, 1]$ that is a strictly proper score function for Γ .

We've seen WSM \Rightarrow properties, the hard direction is the other way.

Proof of Theorem 3: Their Three-Step Argument

Step 1: Construct s

Key property used:
sybilproofness.

Sybilproofness constrains how $f(r, m, M, \omega)$ can depend on m and M separately. By considering groups of clones with equal wagers, the authors show that f must collapse to a single function:

$$f(r, m, M, \omega) \mapsto s(r, \omega, M)$$

Step 2: Extend to Arbitrary m

Properties: sybilproofness + density of the rationals.

For $m_i = a_i/2^k$, replace player i by a_i clones each wagering $1/2^k$. Apply Step 1; sybilproof gives player i 's payout as the WSM formula.

Effect: WSM formula holds for all possible wager vectors.

Step 3: Force $s \in [0, 1]$

Properties: non-negativity of payouts.

Two-player game: player 1 wagers ϵ , player 2 wagers $M - \epsilon$. Non-negativity of $\Pi_1 \geq 0$ forces $1 - (M - \epsilon)s(r_2, \omega, M)/M \geq 0$. Taking $\epsilon \rightarrow 0$: $s \leq 1$.

Truthfulness forces \tilde{s} strictly proper for Γ .

Effect: $\tilde{s} \in [0, 1]$ and is strictly proper. \square

The Incentive Calibration Problem

A fundamental practical challenge with WSM. Payoffs are driven by differences between individual scores and the weighted average:

$$\Pi_i - m_i = m_i(s(r_i, \omega) - \bar{s}).$$

If all players are well-informed and report similar values, the scores $s(r_i, \omega)$ cluster close together and actual money transfers are tiny relative to wagers.

Example: Predicting NBA Scores. Four analysts predict the final score of Lakers vs. Warriors. Outcome $\omega \in [80, 140]$, score $s(r, \omega) = 1 - (\omega - r)^2 / (60)^2$. Assume $\omega = 120$.

- Two experts guessed 115 and 125; they're off by ± 5 points: $s = 1 - 25/3600 \approx 0.993$.
- Two casual fans guessed 100 and 140; they're off by ± 20 points:
 $s = 1 - 400/3600 \approx 0.889$.
- The difference is only 0.104. On a \$100 wager this is only a \$10 differential.

The Core Issue

The score function spans 60 points but all well-informed forecasters report in $[115, 125]$. We need a mechanism that *adapts* to focus sensitivity where the informed reports actually land.

Extension §6.1: Adaptive Weighted-Score Mechanism

The key idea. Let the score function's effective range be informed by what the other forecasters are actually reporting, rather than the original prediction market's bound.

Formal construction. With n players, let \mathcal{S} be all subsets of size $\lceil n/2 \rceil$ (split into sub-games). For player i , let $\mathcal{S}_i = \{G \in \mathcal{S} : i \in G\}$ and $m'_i = m_i/|\mathcal{S}_i|$ (equal wagers to each subgame). For each sub-game $G \in \mathcal{S}_i$:

- Players not playing in G are the **calibrators** — they reveal reports first, forming a reference set $R_G = \{r_\ell : \ell \in \mathcal{N} \setminus G\}$.
- The score function s^{R_G} is parameterised by R_G , giving maximum sensitivity over the calibrators' report range.
- Players in G (the forecasters) compete in a mini-WSM using s^{R_G} .
- Since $i \in G$ cannot influence s^{R_G} (they are not in $\mathcal{N} \setminus G$), their report does not affect the scoring function they face. **Truthfulness is preserved.**

Back to the NBA Example

The 4 analysts split into calibrators and forecasters. If calibrators report in $[115, 125]$, the score function for forecasters is set to have maximum sensitivity over $[115, 125]$ rather than $[80, 140]$.

Extension §6.1: Adaptive WSM — Properties and Intuition

Adaptive WSM Properties

Property

Budget Balance	✓
Anonymity	✓
Truthfulness	✓
Normality	✗
Sybilproofness	✗
Individual Rationality	✓
Monotonicity	✓

Idea Behind Normality and Sybilproofness Failing

Normality fails: Player i 's report enters the calibration set of sub-games where i is a calibrator (not a competitor). Changing i 's report changes the scoring function faced by forecasters in those games, breaking the zero-sum relationship between players' payoffs.

Sybilproofness fails: If expert 1 splits into E_1, E_2 , the two accounts land in different sub-games with different calibration sets. E_1 's calibration may be influenced by E_2 's report, changing the total payout in a way a single expert could not achieve.

Summary and Key Takeaways

What This Paper Achieved

- 1 Presents a general space of mechanisms; **wagering mechanisms**.
- 2 Discusses seven **naturally desirable properties** for **self-financed**, honest forecasting.
- 3 Constructs the **weighted-score mechanism**, a generalization of previous work, satisfying all seven.
- 4 Proves WSM is **unique**: five properties \Leftrightarrow WSM.
- 5 Explores an extension for better calibration.